HOW TO
WORLD'S FIRST ARM BASED HPC CLUSTER

VIEWPOINT
WHY FULL FLASH STORAGE IS GAINING MOMENTUM

NEWSFEED
VMWARE STRATEGY CLOUD NATIVE APPS

HPC REVIEW
THE REFERENCE MEDIA IN HIGH-PERFORMANCE IT SOLUTIONS by HPC TODAY
www.hpcreview.com

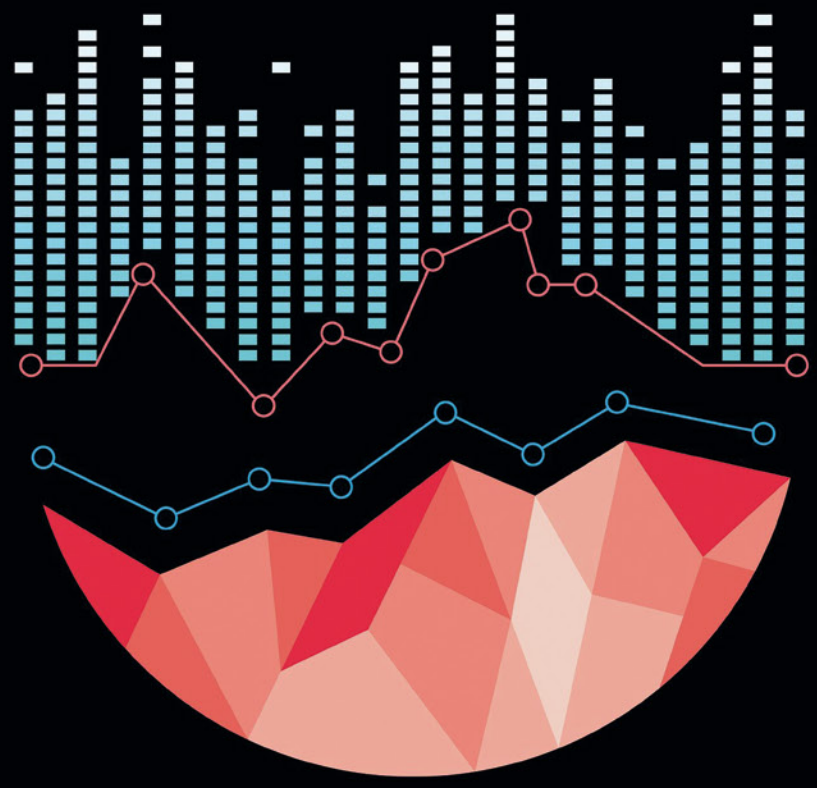

SOFTWARE DEFINED X

POWERING TOMORROW'S DATACENTERS

#9 GLOBAL EDITION

GIGABYTE™

A LEADING GPU SERVER, AT YOUR FINGERTIPS

G190-H44 ***GPU Server***

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

Available at:

MICROTRONICA
A DIVISION OF ARROW

b2b.gigabyte.com

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

FROM HARDWARE TO SOFTWARE: DATACENTERS REDEFINED

WELCOME TO THE NINTH ISSUE OF HPC REVIEW! In this issue we look closely at two strong trends in IT: virtualization and hyperconvergence. After some years of definition, proof of concept and operational validation, both technologies are now mature enough to go invest the datacenters as well as businesses, academic and industry organizations that are becoming more closely linked to them. This strategic link has gained considerable attention since the cloud and connectors applications have established bridges between businesses and data centers. Resilience, redundancy, security, authentication have always distinguished them from enterprise infrastructures. It is now the transition to software X defined that characterizes them, to gain flexibility, agility and competitiveness. Welcome to the dawn of a new world governed by the software.

HAPPY READING!

CONTENTS

HPC REVIEW

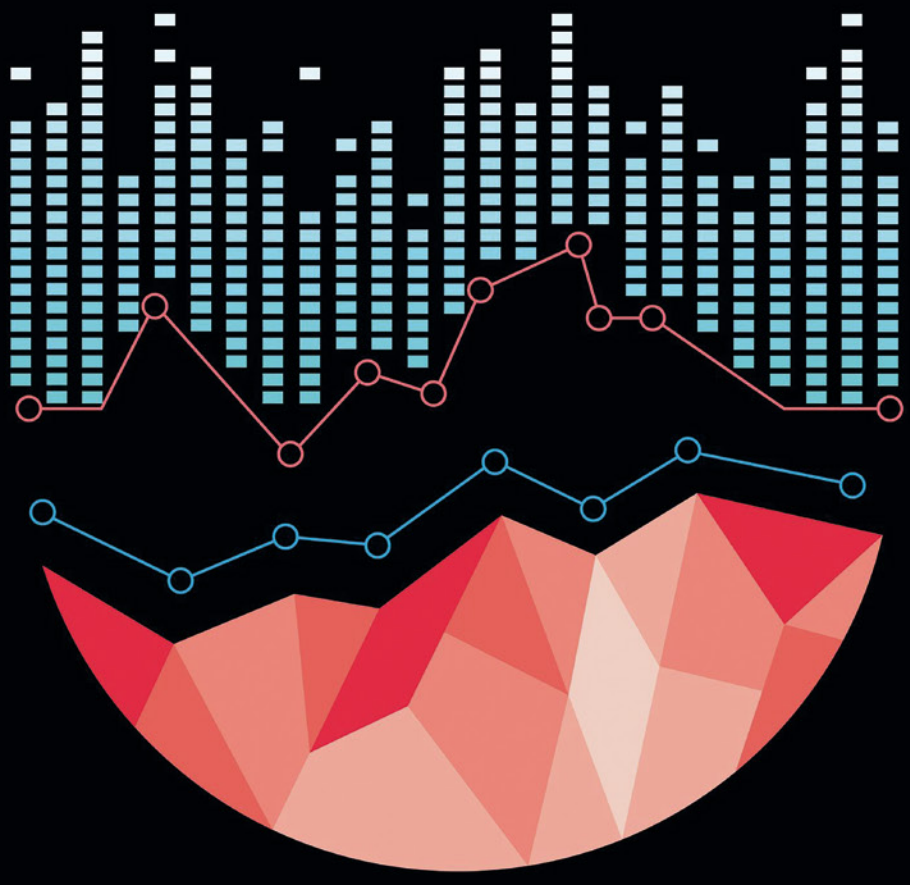

SOFTWARE DEFINED X

POWERING TOMORROW'S DATACENTERS

ASUS®
IN SEARCH OF INCREDIBLE

intel

ESC4000 G3

Green500 Champion Successor

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

HETEROGENEOUS COMPUTING: A NEW PARADIGM FOR THE EXASCALE ERA

STEVE CONWAY
RESEARCH VP, IDC HIGH PERFORMANCE COMPUTING GROUP

The worldwide high-performance computing (HPC) market is already more than seven years into the petascale era (June 2008–present) and is looking to make the thousandfold leap into the exascale era before the end of this decade. This pursuit is global in scope. IDC expects the United States, the European Union, Japan, China, and Russia to vie with each other to reap exascale computing's anticipated substantial benefits for scientific advancement, industrial-economic competitiveness, and the quality of human life.

But as many HPC experts have noted, achieving reasonable exascale performance in this compressed time frame presents an array of daunting challenges that cannot be met only through evolutionary extrapolations from existing technologies and approaches. These challenges include, but are not limited to, the following:

SYSTEM COSTS (FLOPS/DOLLAR). Twenty years ago, the world's leading HPC sites spent $25 million to $30 million for the most powerful supercomputers available. Today's single-digit petaflops supercomputers often cost over $100 million. Early exaflop systems could cost $500 million to $1 billion each. This cost escalation will be difficult to sustain. Anything that can increase the flops/dollar ratio will be welcome.

APPLICATION PERFORMANCE (TIME/SOLUTION). This perennial challenge grows continually as HPC users seek to scale their applications to new, larger systems. With clock rates stalled, future performance gains must come almost entirely from increased parallelism, resulting in tremendous concurrency requirements for exascale computing. A 1GHz machine would need to perform a billion independent operations every clock tick. Over time, many large science problems will be able

DURING THE PAST DECADE, CLUSTERS LEVERAGING THE ECONOMIES OF SCALE OF X86 PROCESSORS BECAME THE DOMINANT SPECIES OF HPC SYSTEMS — DOUBLING THE SIZE OF THE GLOBAL HPC SERVER MARKET FROM ABOUT $5 BILLION IN THE EARLY 2000S TO $9.5 BILLION IN 2010.

to scale to this level. Other problems will lack the required concurrency for single runs but may make use of extreme-scale systems to run ensemble calculations. Automotive design engineers, for example, have greatly increased the number of parametric runs — along with the resolution of each run — that can occur in their allotted phase of the design cycle.

SPACE AND COMPUTE DENSITY REQUIREMENTS (FLOPS/SQUARE FOOT). A worldwide IDC study revealed that most HPC sites are struggling mightily with datacenter space limitations. Two-thirds of the sites were planning to expand or build new HPC datacenters. Half of the sites planned, or had already begun, to distribute their HPC resources to multiple locations.

ENERGY COSTS FOR COMPUTATION AND DATA MOVEMENT (FLOPS/WATT, BYTES/WATT). Last but not least, power has become both a significant design constraint and a major contributor to cost of ownership. With voltage scaling slowing dramatically, power is no longer holding constant as we grow the transistor count with Moore's law, resulting in processor designs that are power constrained today and becoming more so with each new IC generation. Performance in this era is determined largely by power efficiency, so the great challenge in system design is making processors and data movement more energy efficient without overly compromising performance. The rapid growth in HPC system sizes has elevated energy requirements. Today's largest HPC datacenters consume as much electricity as a small city, and multi-petascale and exascale datacenters promise to devour even more. Energy prices have risen substantially above historic levels, although prices have moderated from their 2008 highs. Another element in this «perfect storm» is that HPC datacenter power and cooling developments are occurring at a time of growing sensitivity toward carbon footprints and global climate change. Finally, some of the biggest HPC datacenters worry that their local power companies may balk at fully supplying their future demands. One such site, already seeing the need for a 250-megawatt datacenter, may have to go off the power grid and build a small nuclear reactor.

THE HETEROGENEOUS COMPUTING PARADIGM

During the past decade, clusters leveraging the economies of scale of x86 processors became the dominant species of HPC systems — doubling the size of the global HPC server market from about $5 billion in the early 2000s to $9.5 billion in 2010. The reigning paradigm has been to advance peak performance by deploying larger and larger clusters containing more and more standard x86 CPU cores.

But x86 processors were never designed to handle all HPC applications well, and x86 single- threaded performance started to hit a heat and power wall half a dozen years ago. It is becoming increasingly clear that although x86 processor road maps lay out substantial advances, the paradigm of sole dependency on x86 processors will not suffice to meet the challenges associated with achieving exascale computing in this decade.

In recent years, an alternative paradigm, «heterogeneous computing» has gained market momentum for addressing these challenges. This emerging paradigm augments x86 CPUs with accelerators, primarily GPGPUs (henceforth to be called GPUs), so that each processor type can do what it does best. GPUs are particularly adept at handling the subs-

tantial number of codes, and portions of codes, that exhibit strong data or thread-level parallelism. That makes GPUs the heirs apparent to vector processors, except that GPUs benefit from far greater economies of scale and related competitive advantages. IDC research shows that the worldwide PC market for discrete graphics processing units alone was worth about $4 billion in 2010.

The heterogeneous computing paradigm is ramping up nicely across the HPC market as a whole. IDC's 2008 worldwide study on HPC processors revealed that 9% of HPC sites were using some form of accelerator technology alongside CPUs in their installed systems. Fast-forward to the 2010 version of the same global study and the scene has changed considerably. Accelerator technology has gone forth and multiplied. By this time, 28% of the HPC sites were using accelerator technology — a threefold increase from two years earlier — and nearly all of these accelerators were GPUs. Although GPUs represent only about 5% of the processor counts in heterogeneous systems, their numbers are growing rapidly.

Heterogeneous computing is making its greatest impact at the high end of the HPC market. GPUs first appeared on the TOP500 list of the world's supercomputer sites (www.top500.org) in 2008. By June 2011, three of the top 10 systems on the list employed GPUs. And in October 2011, the U.S. Department of Energy's Oak Ridge National Laboratory unveiled plans to upgrade the number one

U.S. supercomputer to a successor system («Titan») with a planned peak performance of 20–30 petaflops by complementing more than 18,000 x86 CPUs with an equal number of GPUs. Following that, the Texas Advanced Computing Center revealed plans for a heterogeneous supercomputer («Stampede») that initially targeted 10 peak petaflops by combining two petaflops of x86 CPUs with eight petaflops of MIC accelerator processors.

The adoption of heterogeneous computing by these and other bellwether HPC sites indicates that GPUs are moving out of the experimental phase and into the phase where they will be increasingly entrusted with appropriate production-oriented, mission-critical work.

DEFINITIONS

CLUSTER: IDC defines a cluster as a set of independent computers combined into a unified system through systems software and networking technologies. Thus, clusters are not based on new architectural concepts so much as new systems integration strategies.

HETEROGENEOUS PROCESSING: Heterogeneous processing and the synonymous term heterogeneous computing refer to the use of multiple types of processors, typically CPUs in combination with GPUs or other accelerators, within the same HPC system.

HIGH-PERFORMANCE COMPUTING: IDC uses the term high-performance computing to refer to all technical computing servers and clusters used to solve problems that are computationally intensive or data intensive. The term also refers to the market for these systems and the activities within this market. It includes technical servers but excludes desktop computers used for technical computing.

HETEROGENEOUS COMPUTING BENEFITS FOR THE EXASCALE ERA

The benefits of the heterogeneous computing paradigm for HPC are interrelated and address some of the most important exascale challenges:

SYSTEM COSTS. GPUs and related accelerators can provide a lot of peak and Linpack flops for the money. Especially for HPC sites seriously pursuing the upper ranges of the TOP500 supercomputers list, bolting on GPUs can provide a kind of flops warp drive, rocketing Linpack performance to where almost no one has gone before. Witness China's Tianhe-1A supercomputer, which supplemented x86 processors with GPUs to seize the number one spot on the November 2010 TOP500 list. To achieve this feat, Tianhe-1A employed 14,336 x86 CPUs and 7,168 GPUs. NVIDIA suggested at the time that

it would have taken «50,000 CPUs and twice as much floor space to deliver the same performance using CPUs alone.» In any case, by June 2011, three of the top five systems on the list employed GPUs. And in October 2011, as noted earlier, the U.S. Department of Energy's Oak Ridge National Laboratory unveiled plans to upgrade the number one U.S. supercomputer to peak performance of 20–30 petaflops by complementing more than 18,000 x86 CPUs with an equal number of GPUs.

SPEED. HPC users have reported excellent speedups on GPUs, frequently in the 3–10x range, especially for codes, or portions of codes, that exhibit strong data parallelism. GPUs are already enabling real-world advances in HPC domains, especially the life sciences, financial services, oil and gas, product design and manufacturing domains, and digital content creation and distribution. GPUs are a particularly promising fit for molecular dynamics simulations, which extend across multiple applications domains.

SPACE AND COMPUTE DENSITY. At a time when many HPC datacenters are approaching the limits of their power and space envelopes, GPUs promise to deliver very high peak compute density. A contemporary GPU may contain as many as 512 compute cores, compared with 4 to 16 cores for a contemporary CPU. Keep in mind, however, that heterogeneous computing is heterogeneous for a reason: Each processor type, CPU, and accelerator is best at tackling a different portion of the problem-solving spectrum.

ENERGY COSTS. The rapid growth in HPC system sizes has caused energy requirements to skyrocket. Today's largest HPC datacenters consume as much electricity as a small city, and exascale datacenters promise to devour even more — an estimated 120 megawatts or more, if they were implemented using existing technologies. The Department of Energy's exascale goal is to bring that number down to no more than 20 megawatts for an exascale system deployment. This is desired to avoid massive increases in energy costs, to ensure the availability of adequate energy supplies from local utilities, and to keep datacenter spatial requirements within reason. GPUs can be valuable partners with CPUs in heterogeneous computing configurations by significantly improving energy efficiency on the substantial subset of codes (and portions of codes) exhibiting strong data parallelism.

ADOPTION BARRIERS

As a relatively new technology — at least as used for computing — GPUs have encountered adoption barriers that IDC expects to ease over time. HPC buyers report the following main barriers to more extensive GPU deployment:

EASE OF PROGRAMMING. Despite the availability of useful tools such as CUDA, OpenCL, and The Portland Group's directives-based compiler that's designed to transform Fortran or C source code into GPU-accelerated code, HPC buyers and end users generally report that programming GPUs remains more challenging than the more familiar approaches to programming standard x86 processors. This barrier will likely continue to drop over time as familiarity with GPU programming methods grows — through the more than 450 universities offering GPU curricula today and as the GPU programming methods advance.

MEDIATED COMMUNICATION. Another issue frequently cited by HPC users is the fact that GPUs today are typically implemented as co-processors that need to communicate with x86 or other base processors via PCI Express channels that are comparatively slow — at least when weighed against implementing the CPU and GPU on the same die. This mediated communication affects some applications more than others. It has not prevented HPC users from achieving impressive time-to-solution speedups on a growing number of application codes.

WAITING FOR FUTURE CPU GENERATIONS. Some HPC users believe that waiting to see

what improvements future-generation x86 processors deliver is a risk worth taking, compared with the effort of learning how to program GPUs and adapting portions of their codes to run on GPUs. And because GPUs are still relatively new devices for high-performance computing, some users worry that the substantial effort to rewrite their codes could be wasted if GPU architectures evolve in a new direction or if GPUs are not an enduring phenomenon in the HPC market. This wait-and-see group has been declining as GPUs have increased their influence in the global HPC market and as directive-based programming of GPUs has become more prevalent.

TRENDS

Heterogeneous computing, which today typically couples x86 processors with GPUs implemented as coprocessors, is an important new paradigm that is increasingly taking its place alongside the existing paradigm of pure-play x86-based HPC systems.

An important sign of GPUs' momentum is the spread of GPU-related academic offerings. NVIDIA, which supplies educational materials for parallel programming, reports that its CUDA parallel programming language is being taught at 478 universities in 57 countries. The list includes MIT, Harvard, Stanford, Cambridge, Oxford, the Indian Institutes of Technology, National Taiwan University, and the Chinese Academy of Sciences.

For reasons stated earlier (refer back to the Benefits section), heterogeneous computing is proving especially attractive to large HPC sites that are pushing up against the boundaries of computational science and engineering, as well as energy and spatial boundaries. Hence, heterogeneous computing looks particularly attractive as a new paradigm for the exascale computing era that will begin later in this decade. Heterogeneous computing involving GPUs is also making inroads into smaller research sites and industrial organizations.

Keep in mind, however, that x86 processor technology is not standing still and promises to remain the HPC revenue leader through 2015. In addition, accelerator technology will be available from an increasing number of vendors and in a growing array of «flavors,» giving users more options.

CONCLUSION

Heterogeneous computing that today typically couples x86-based processors with GPUs implemented as coprocessors is an important emerging paradigm in the worldwide HPC market — especially as a strategy to help meet the daunting challenges of the coming exascale computing era. IDC believes that heterogeneous computing will be indispensable for achieving exascale computing in this decade.

GPUs have rapidly emerged from the experimental phase and are used today for production- oriented tasks such as seismic processing, biochemistry simulations, weather and climate modeling, computational finance, computational fluid dynamics, and data analysis. They have tripled their worldwide footprint at HPC sites in the past two years alone, they have become more indispensable for attaining prominence on the closely watched TOP500 supercomputers lists, and they have enabled a growing number and variety of real-world achievements.

The adoption of heterogeneous processing involving GPUs by some of the world's leading HPC sites indicates that this paradigm is moving beyond the experimental phase, and GPUs are increasingly being entrusted with appropriate portions of production-oriented, mission-critical workloads.

As GPU hardware and software technologies advance, as more university students and others learn how to exploit GPUs, and as more GPUs become available to the world's most creative scientific, engineering, and computational minds, IDC believes GPUs will play an increasingly important role in the global HPC market, complementing x86 processors within the HPC ecosystem.

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 Infrastucture

Customer Relationship Management (CRM)

Online Archiving

Video Editing/ Photo Sharing

Enterprise Content Management (ECM)

Virtual Tape Library (VTL)

OCZ STORAGE SOLUTIONS | A Toshiba Group Company

BECHTLE bechtle.co.uk | CCL cclonline.com | SPAN.COM span.com

ocz.com

HPC REVIEW

VMWORLD 2015 EUROPE

CLOUD NATIVE APPS ARE THE FUTURE ACCORDING TO VMWARE

DURING ITS OPENING KENOTE SESSION AT VMWORLD EUROPE 2015, CEO PAT GELSINGER HIGHLIGHTED THE THREE STRATEGIC ASPECTS FOR ITS BUSINESS. THESE ARE THE TECHNOLOGIES AROUND THE DATA CENTER, THE SOFTWARE STRATEGY FOR BUSINESS MOBILITY, AND THE IMPACT OF DIGITAL TRANSFORMATION IN THE BUSINESS WORLD.

THE GOAL IS TO CHANGE THE MODEL NETWORK SECURITY END TO END BY INTEGRATING INNOVATION FOR APPLICATIONS AND DEVICES AND DATA.

Gelsinger further stressed that VMware would continue to support the digital business transformation. The motto of VMworld this year was «Ready for any» with a subtitle that encompasses the three aforementioned topics: «One cloud, Any Application, Any device». The justification for this is that strong axis applications are doing to change the world to digital and thereby our way of working. VMware's goal is to provide a single infrastructure layer to help businesses to support their application traditional as well as «Cloud native. « The challenge is to manage applications and workloads as evenly as possible. Another highlight is to offer companies the ability to access applications from any mobile terminal in fluid mode without disconnection. With another hot topic regarding the underlying security. The goal is to change the model network security end to end by integrating innovation for applications and devices and data. Gelsinger then described the five fundamental axes of engagement required for a company to succeed in a digital world.

1 Asymmetry in Business: elephants must learn to dance

Companies must learn to innovate as a start-up, and deliver as a business. For Gelsinger, the challenge is longer to go higher, further and stronger, but to be more reactive, faster and

more responsive. Companies that will prevail in this new era are no longer the biggest or the oldest, This is no longer the company greater or old, but the fastest.

2 Entering the professional era of Cloud: unified Cloud is the future

Business infrastructures will become increasingly hybrid with the advent of connected objects, distributed infrastructure and services and applications heterogeneity. This is the new challenge for the CIO and the major issue around cloud technologies and unification of resources.

3 The Security Challenge: Protecting people, apps and data

This new era also poses major challenges around safety. Overall, IT budgets are in a stabilization phase, and the only growing budgets are those related to security. The budgets to improve security are becoming increasingly important in order to compensate for environments that are not necessarily secured enough to begin with. It has therefore become imperative to rethink security as a whole and to use new technologies around the Network like SDN (Software Defined Network). For it is a necessary step to support the infrastructure as it evolves. For Gelsinger, we are entering a new era for security.

4 The next wave of innovation: proactive technology

For Gelsinger, proactive technology must help to automate almost everything. The capacity to capture and analyze data in time real must lead to intelligent systems who can analyze human behavior and adapt itself. With strong artificial intelligence implications around the manipulation, treatment and security around the data.

5 Tech-driven exchange reshapes the S & P 500: Taking Risks is the lowest risk

Gelsinger estimated that 40% of businesses no longer exist in 10 years by not taking technological change in time to propose new services. It is therefore more risky not to take risks to take to address the new corners technology.

3 BILLION PEOPLE CONNECTED IN 2025

Gelsinger announced that in 2025, 3 billion people will be connected, with the cloud becoming the foundation of this digital world. And one huge issue around security, which it becomes one of the basic prerequisites. At the heart of these digital services, Gelsinger firmly believes that companies need to transform themselves to succeed and survive. He believes that VMware and its parent company, EMC, need to improve integration to be relevant to a changing market, including in terms of architecture with EMC storage technologies and components. As a sign of things to come, the revenue model of VMware has evolved significantly to reflect these changes. VMware's virtualization revenues have declined down between 35 and 40% of its turnover. The rest proceeds from its non-virtualization activities. End User Computing activities are said to be evolving significantly, however details have not yet been revealed.

FEDERATION : AN INTEGRATED APPROACH

This integrated federation approach is achieved through vCloud Air, with two deals around storage with world-class Cloud storage vendors based on Google and EMC for

Premium Storage. Altogether making part of the EHC initative, Enterprise Hybrid Cloud, and having in common vSphere as a virtualized infrastructure layer. VMware therefore positions itself as a cloud platform technology provider with its VMware platform Unified Hybrid Cloud platform around two strong axes. On the one hand simplifying supply for enterprises with the integration of EVO SDDC Manager vRealize Suite, NSX 6.2, Virtual SAN vSphere 6.1 and 6, and secondly its extension through vCloud Air hybrid networking services, content synchronization mechanisms and workload movement between heterogeneous Cloud, illustrated by a Tech Preview of Skyscraper project.

SERVICE AUTOMATION BETWEEN CLOUDS

VMware also continues to work on automation services for Companies such as data centers. A concrete example: establishing a mirror mode link for synchronization templates between virtual machines in private and public clouds. Or the addition of network extensions between private cloud and public cloud through mechanisms dedicated to long-distance migration through vCenter vMotion. This platform also allows Information collection on the deployed VMs according to the available Cloud load and availability. The addition of a local frontend with remote data allows organizations to better take advantage of use Public Cloud and limit the shadow IT.

DISASTER RECOVERY OFFERS

VMware has also announced new dedicated vCloud Air storage and disaster recovery services to be operational early next year through vCloud Air Disaster recovery services Air vCloud mobile backend-as-a service, vCloud Air Storage object, and finally vCloud Air SQL. The stated aim is to bring to customers a public mode service extension for four types of use: site or data protection, on demand use (e.g. DirecTV) mobile applications, and development of test applications running a public cloud. The service catalog will continue to evolve and improve. Other technologies announced include vSphere integrated containers. This offer integrates technologies vRealize, vSphere, and VSAN NSX, and adds others specifically developed like Instant Clone as the project Bonneville (application container) or Photon OS, a lightweight Linux operating system weighing just 27 megabytes.

A NEW ADDITIONAL FOCUS : BUSINESS MOBILITY

VMware characterized these cloud applications native as being able to go beyond the usual scope assigned to them, through distributed applications in fragmented form, ie microservices, for mobile professionals. This new paradigm puts the user at the center of management and administration proposal through a single sign on gateway. The management of these users being devolved to new VMware Identity Manager Advanced. The highlighted benefits are threefold: increasing the mobility of staff, streamlining business processes and reduce operating costs.

MICROSOFT AND VMWARE PARTNERS FOR MOBILE APP DELIVERY

Also in terms of mobility, the company presented the A² project with its competitor and partner for the occasion, Microsoft. This project has been demonstrated as a Tech Preview for Appvolume, a joint work for AirWatch application management and delivery. The objective is here to offer a framework for provisioning and deplying pplication for Windows 10 users. An initiative coherent with Microsoft's Universal Apps which represents the convergence between the desktop and mobile mode of app delivery. Based on AirWatch, this completes Micvrosoft's offer to cover iOS and Android Smartphones and tablets.

The projected revenue of enterprise Cloud business should skyrocket from $ 22 billion today to 77 billion in 2030. Adoption of hybrid cloud is therefore considered as an industry-wide issue. **JOSCELYN FLORES**

SYLVAIN CAZARD

CEO of VMware France

"WE WORKED TO EXPAND OUR PRODUCT PORTFOLIO BY INTEGRATING THE TECHNOLOGIES NEEDED TO BETTER PREPARE FUTURE BUSINESS NEEDS."

Sylvain Cazard, CEO of VMware France, gives us an update on the status of the strategy of his company shortly after VMworld 2015's announcements draw a wider perimeter.

Can you update us about your company's strategy?

VMware is doing very well in a rather sluggish industry. We are pursuing the same strategy initiated three years ago by our CEO Pat Gelsinger, who has several objectives: first, help transform the datacenters and software-defined datacenters into digital software factories, which are intended to benefit both the infrastructure of companies and their cloud service providers. We also work to improve the future workplace through our Workspace application portal that is constantly evolving to reflect changes in business and mobility needs. As demonstrated by the announcements made at our last VMworld 2015, we have worked to expand our product portfolio by integrating the technologies needed to better prepare for the future needs of enterprises, including the means to assist them to better achieve their digital transformation.

Which of your products portfolio have been improved most significantly?

We have continued to invest in the same direction, both through internal development as well as the needed skills and activities to complement our offer. We also worked on the integration of new components and technologies to support our strategy and continue to anticipate the changing business requirements of our customers. This integration of our technological bricks has one objective: to help them increase their value. Mobility is a growing part of our business. Management of mobile fleet management and of distributed applications while maintaining a strong end to end security are strongly increasing imperatives.

How have your activities evolved?

We of course continue to work on developments at the heart of our business around virtualization. Our customers continue to show their confidence in our offers that highlight the value proposition. It is interesting that it now exceeds traditional virtualization framework that is heart of our business, to extend our other offers and services and take the form of support over time. We therefore speak only more hypervisor, but the integration of a set of software solutions in the service of their activity and their occupations.

Beyond the role of technology provider, our business has also evolved towards more support for enterprises. We rely on it for 200 partners in France, a figure doubled in a year. And an interesting fact: the number of SMEs that gradually tilt in one of our solutions is significant and growing. A strong trend with an interesting corollary around a growing num-

ber of cloud service providers adopt and integrate our technologies.

How do you see the evolution of your business offerings?

Our offer is strengthened particularly in terms of converged and hyperconverged platforms, which further strengthens the coherence of our strategy. Our corporate EVO Rail appliances offer is complemented by a range EVO SDDC (ex EVO Rack), which completes the functional scope with the layer software defined network including NSX. We are accompanied on these initiatives by our eight OEMs: Dell, EMC, Fujitsu, Hitachi Data Systems, Inspur, Netapp, NetOne and SuperMicro. Keep in mind that in order to ensure simplified deployment, the EVO Rail appliances go through a rigorous and complex industrial process of certification and integration. In terms of the typology of any company adopting these two product lines EVO Rail is adopted by SMEs and IT departments of major accounts. Evo SDDC aims for its share of high availability tasks with high performance test and find its place in the larger datacenter.

What is your strategy regarding application containers?

As confirmed by our announcements at VMworld, is a major development for VMware, that we had already identified last year through our partnership with Docker. We continue our developments to provide a comprehensive platform of application containerization with the same guarantees of resilience and availability as the rest of the technical infrastructure that we put in place. Many Tech Previews allowed to lift a veil on the future of application containers at VMware, which is a priority for us. We combine the logic of the portal application a range of services to go towards more automation, as many criteria to allow companies to retain control of their information systems throughout the conversion of the latter, and aim a self-service infrastructure.

How do you see the evolution of business needs in the coming months?

We expect a double evolution. First among our customers that we support for the establishment of their IT transformation plan, adjusted according to their level of maturity. And acceleration of project mode deployments incorporating steps, components, services, and applications involved in the business and take advantage of a growing hybridization. With a decisive progress, based on our ability to accurately calculate the ROI and TCO of the modeled digital transformation projects throughout the cycle of their deployment. Note that Gartner and Accenture have been with us for large projects.

On the other hand, we expect further mergers or consolidation of data centers as they are adopting software defined technology to improve their infrastructure responsiveness to meet their customer's needs. All these developments fit in our strategy and are on our roadmap.

What key elements do you see clear in the near future?

We see Internet of Things with great interest as to support the exponential increase of connected objects, we must continue to improve the resilience and adaptability of the infrastructure that will support the business around of these objects. Other topics that we see with increasing importance in the coming years are OpenStack, hyperconvergence, DevOps and application containerization and of course, software defined network and the related security topics.

INTERVIEW BY **JOSCELYN FLORES**

QUIET, RELIABLE AND EFFICIENT: JOBY AVIATION PAVES THE WAY FOR AN ELECTRIC FUTURE

A rendering of the Joby S2

Electric propulsion is on the verge of causing the biggest changes in aviation since the advent of the jet engine. At first glance, it may seem that the excessive weight (i.e. low specific energy) of today's batteries limits electric aircraft to, at best, a few trivial niches. However, the different properties of electric propulsion compared to traditional combustion power, coupled with recent technology advances, promise to significantly relax typical design constraints for many aircraft configurations, which will allow for new types of aircraft that were previously impractical or impossible. This is particularly true for shorter-range designs, which have traditionally been relatively small and piston-powered.

WHY ELECTRIC PROPULSION ?

Because of the size, weight, and maintenance requirements of piston engines, most piston aircraft designs are limited to a small number of engines (often just one) located in a small number of practical locations. This is why most modern general aviation airplanes and helicopters look very similar to designs from the 1950s. In contrast, electric powertrains are much smaller and lighter, and they are incredibly simple – some having only a single moving part – compared to the relatively ex-

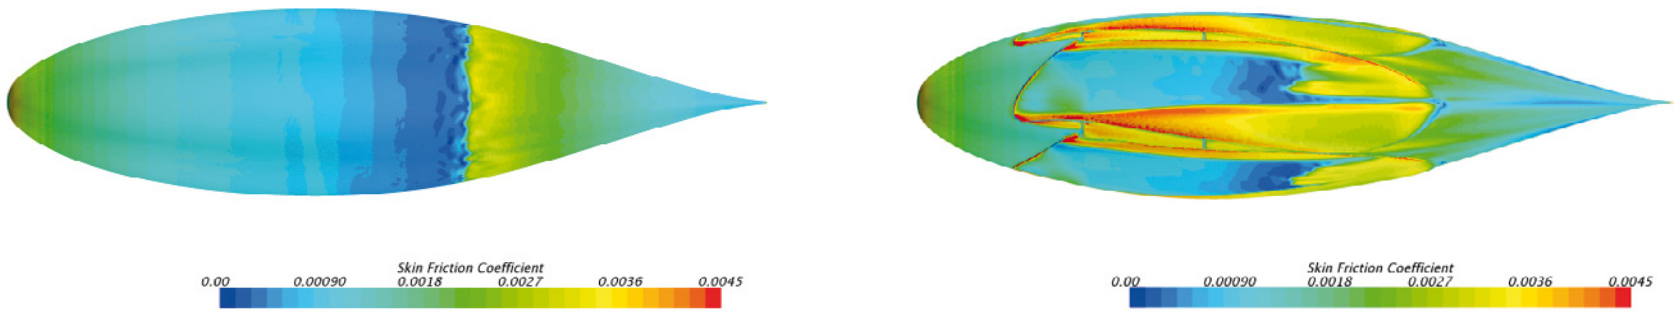

CFD analysis of the S2 nacelles, showing a clean nacelle (left) and a nacelle with folded blades and spinner gaps (right)

treme complexity of piston engines, which include a coolant system, an electrical system, an oil system, a fuel system, and so forth. This reduced complexity translates to much lower maintenance requirements.

While smaller combustion engines suffer from lower power-to-weight and efficiency, electric motors are relatively scale-free. This means that the power-to-weight and efficiency will be similar between, for example, a 1 kW motor and a 1,000 kW motor. An electric powertrain is also about three times as efficient (around 90%-95% compared to around 30%-40%). Electric motors can operate well on a much wider range of RPMs, and they can change RPM relatively quickly.

Electric powertrains are significantly quieter than combustion powertrains, as anyone who has heard an electric car can attest.

LOWER NOISE AND HIGHER EFFICIENCY

While simply replacing a combustion engine with an electric motor will see the benefits of lower noise and higher powertrain efficiency, much greater advantages can be gained by designing an aircraft with electric propulsion in mind from the start. The different properties of electric propulsion mean that aircraft can effectively employ a large number of small motors without incurring an undesirable amount of complexity (and maintenance costs) and without compromising on motor weight or performance. These motors can be located in a much larger range of positions on the aircraft, due to their relatively low weight and small size. Additionally, the drawbacks of carrying motors that are only used in some portions of the flight (e.g. takeoff and landing) are relatively minor, since the motors themselves are so light.

While traditional propulsion installations often compromise aircraft performance – for example, the scrubbing drag caused by a tractor propeller increasing the velocity over the fuselage – the flexibility of electric propulsion allows for propulsion installations that actually result in beneficial aerodynamic interactions. One such example is locating propellers on the wingtips, where they can recapture some of the energy lost to the wingtip vortices.

With its expertise in electric motor design and fabrication, high-fidelity aerodynamic analysis, and composite airframe design and fabrication, Joby Aviation is fully capitalizing on the promise of this new technology to develop several aircraft providing capabilities that were never before possible. However, due to the complex nature of these interactions and the lack of previous designs to extrapolate from, a large amount of high-order aerodynamic analysis must be performed in the design process. For this reason, Joby Aviation has leaned heavily on CFD analyses using STAR-CCM+® in the development of its unconventional designs.

PROJECT 1: THE JOBY S2

Joby Aviation's main development effort is the S2 Vertical Takeoff and Landing (VTOL) aircraft, shown in Figure 1, which addresses the

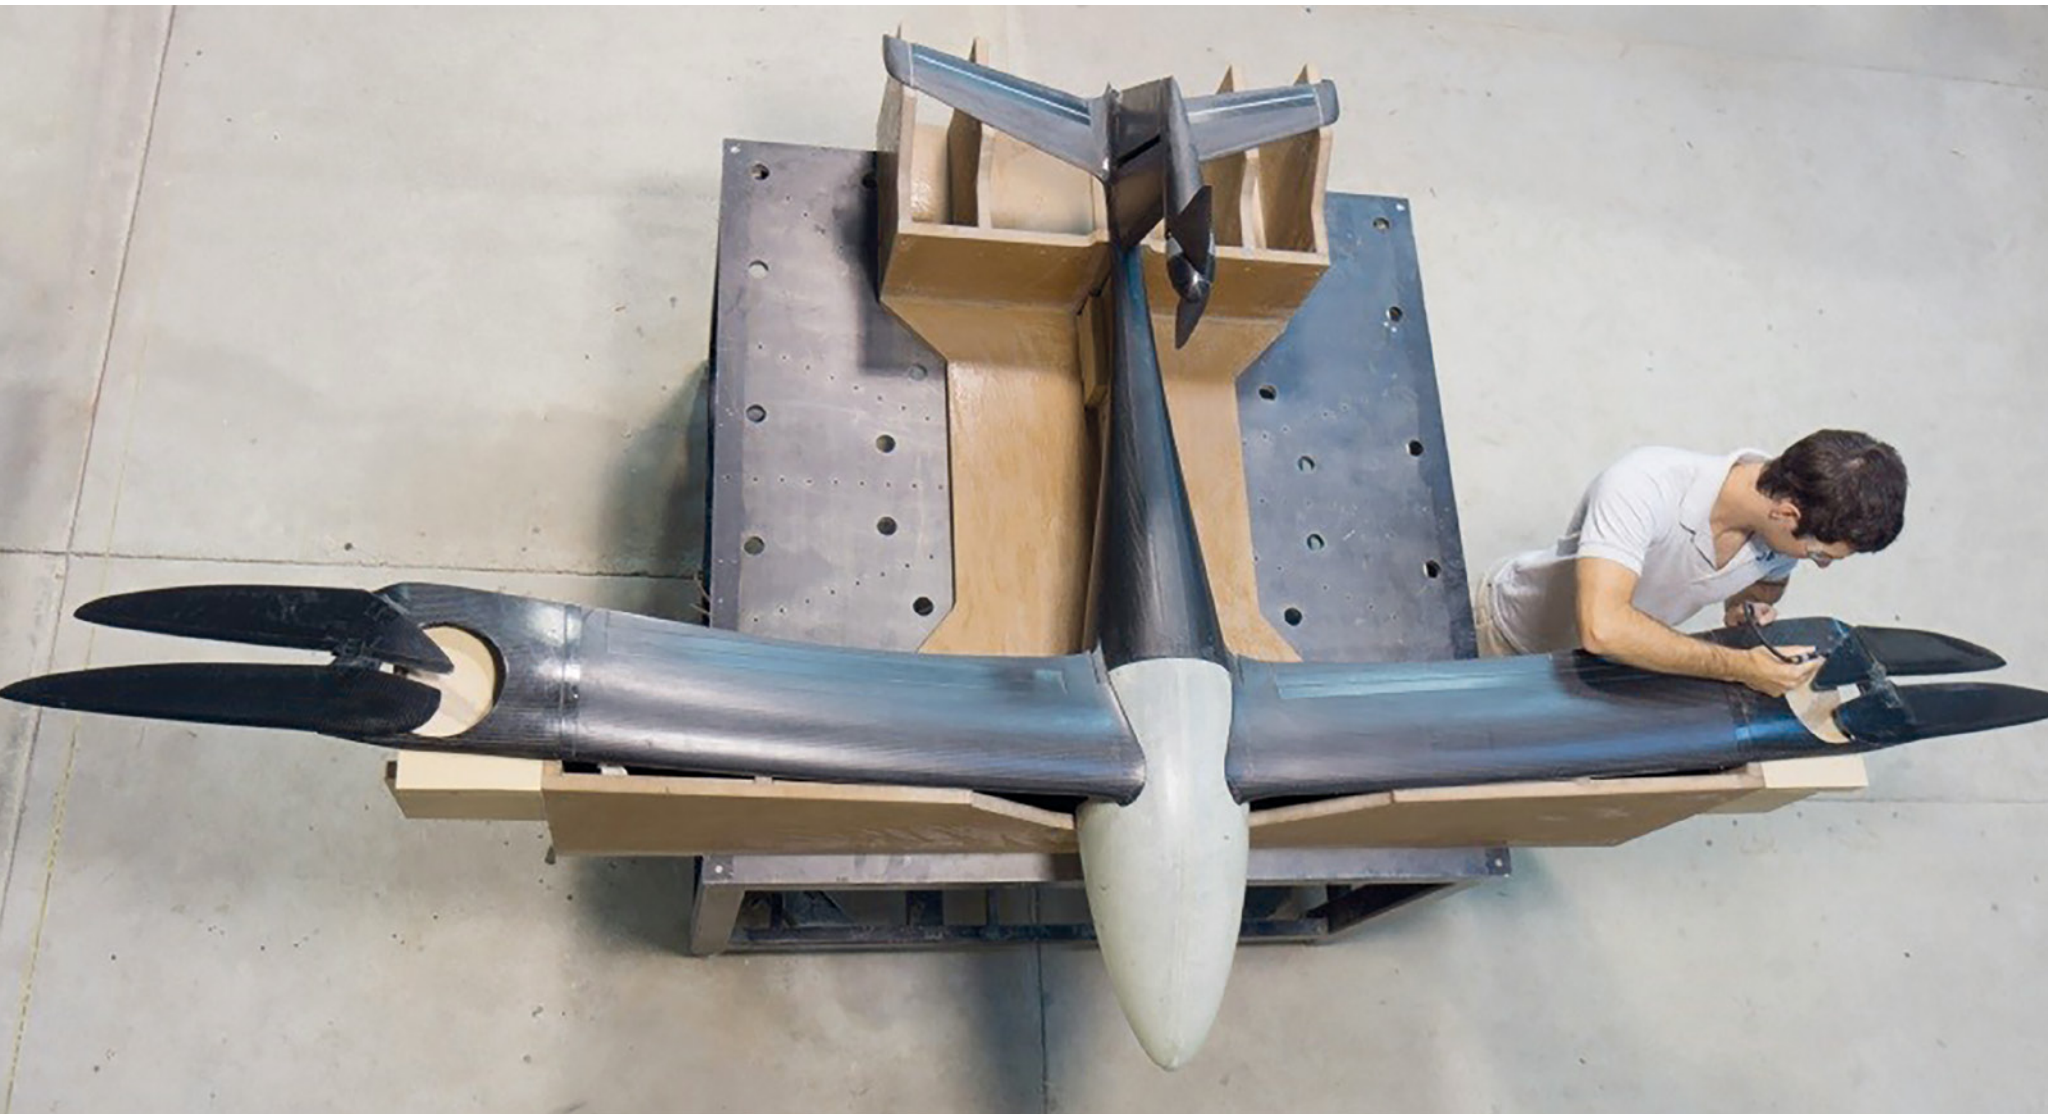

high noise, high operating costs, low speed, and relatively low safety levels that, together, have severely limited the proliferation of conventional VTOL aircraft of this size (helicopters). The S2 employs multiple propellers in takeoff and landing to increase safety through redundancy. In cruise, most of these propellers fold flat against their nacelles to reduce drag. The design of these propeller blades is a compromise between propeller performance and the drag of the nacelles with the blades folded, and higher-order tools were required to properly analyze this tradeoff. A variety of propeller designs were assessed under various operating conditions in STAR-CCM+, and the nacelle was analyzed in the cruise configuration using the γ-Reθ transition model. One such nacelle geometry can be seen in Figure 2, where both the unmodified clean nacelle and the same nacelle with the folded blades and spinner gaps are shown. Such results indicate where reshaping the propeller blades may increase laminar flow and reduce cruise drag.

The Lotus during assembly, in cruise configuration

PROJECT 2 : JOBY LOTUS

Another Joby Aviation project is the Lotus aircraft, shown in Figure 3, which is exploring a novel VTOL configuration on the 55-pound UAV scale. In this aircraft, two-bladed propellers on each wingtip provide thrust for vertical takeoff. After the aircraft picks up enough forward speed for sufficient wing lift, each set of two blades scissors together and the individual blades become wingtip extensions, forming a split wingtip. A tilting tail rotor provides pitch control during takeoff and landing and propels the aircraft in forward flight. The takeoff and cruise configurations of the Lotus are illustrated in Figure 4. As one may expect, the design of these wingtip blades – the span, airfoil choice, twist and chord distribution, pitch, and dihedral – was an interesting compromise between propeller and wingtip performance. Dozens of CFD simulations were run on different combinations of these

CFD analysis of the Lotus in cruise configuration

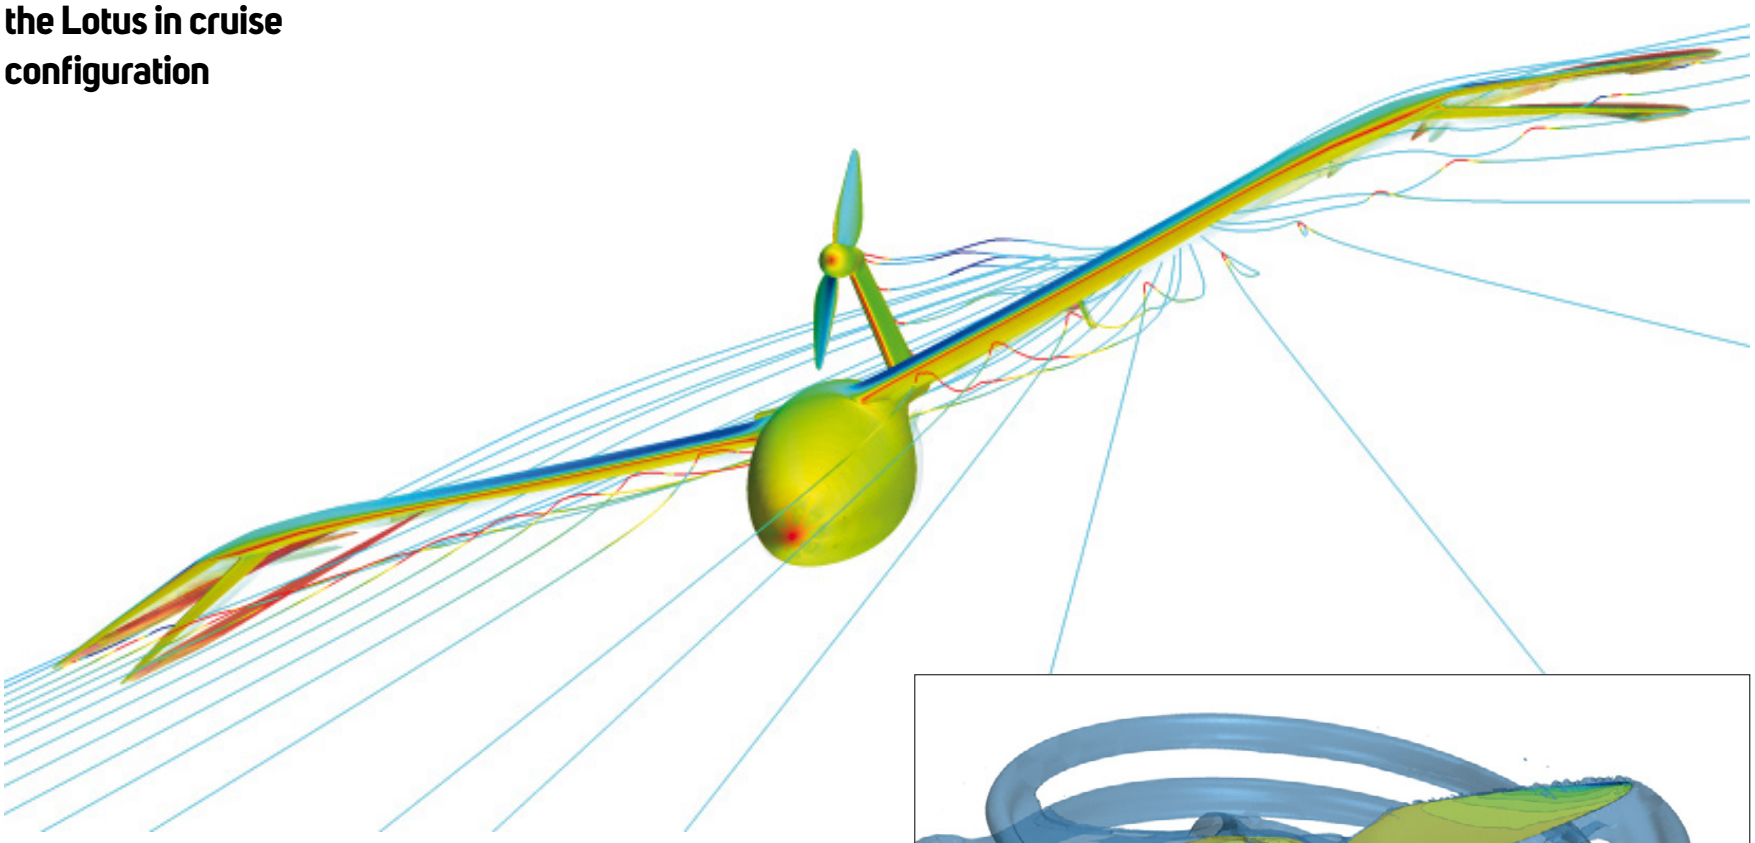

design variables in the cruise configuration, to maximize the cruise performance within the constraints of the configuration. At the same time, the performance of these blades in the propeller configuration was also analyzed with CFD to validate lower-order design methods. Example results from some of these simulations are shown in Figures 5 and 6.

PROJECT 3: LEAPTECH

The third project Joby Aviation is participating in is LEAPTech (Leading Edge Asynchronous Propeller Technology), a partnership with NASA and Empirical Systems Aerospace. The goal of this design is to investigate potential improvements in conventional fixed-wing aircraft through electric propulsion. A row of small propellers is located along the leading edge of the wings and, during takeoff and landing, these propellers increase the velocity (and, therefore, the dynamic pressure) over the wings. This increases the lift produced by the wing and allows for a smaller wing to be used for the same stall speed constraint. Since many small aircraft use a wing sized to meet a stall speed constraint but too large for optimal cruise performance, this smaller wing

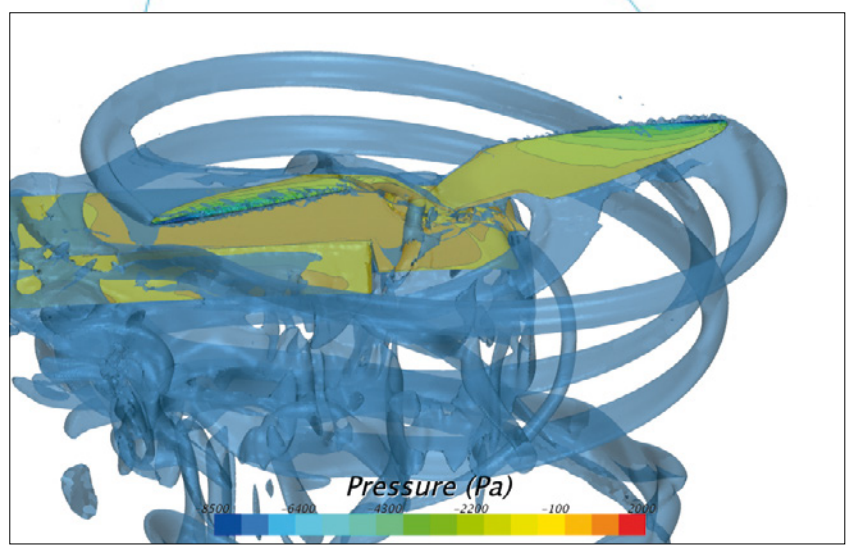

CFD analysis of the Lotus wingtip propeller at takeoff

allows for more efficient cruise. Additionally, the ride quality is significantly improved due to the higher wing loading. However, the performance of this blown wing is difficult to analyze with lower-order tools, particularly since much of the required analysis occurs around stalling conditions. Therefore, a large number of CFD simulations were performed in the design process, looking at various combinations of propeller sizes and powers, wing aspect ratios and sizes, angles of attack, etc. To reduce the computational expense, the propellers were modeled as actuator disks with the body force propeller method in STAR-CCM+, which negated the need to resolve the actual blade geometry, drastically decreasing the required mesh size.

The first phase of testing this configuration was to build the full-scale wing, propel-

lers, and motors, and mount them above a modified semi-truck which was run at takeoff speeds on the runway at NASA Armstrong Flight Research Center. An example CFD solution of this configuration is shown in Figure 7, and the experimental test apparatus is shown in Figure 8.

Outside of takeoff and landing, these leading-edge propellers are planned to fold against their nacelles – similar to the S2 propellers – and wingtip propellers, as mentioned above, will provide propulsion. Although lower-order analysis methods were evaluated for estimating the drag and efficiency impact of operating these propellers concentric with the wingtip vortex, unsteady CFD proved to be the most reliable analysis method. A range of design parameters were analyzed, and one such solution is shown in Figure 9. A flight demonstrator is planned for flights beginning in 2017; a rendering of this aircraft is shown in Figure 10.

CONCLUSION

Joby Aviation is quickly advancing the state of general aviation aircraft with its revolutionary electric propulsion concepts, and simulation is playing a big role in understanding the complex nature of their state-of-the-art ideas and in the design and development of their unconventional systems. The S2, Lotus, and LEAPTech designs show great promise towards an electric future in aviation never before possible.

HPC ACADEMY

TEACHING, PROMOTING AND DEMOCRATIZING SUPERCOMPUTING

CHRISTOPHE RODRIGUES, Executive Director, The HPC Academy

« WE BELIEVE THAT HPC IS NOW OUT OF ITS IVORY TOWER AND THAT, TO PUT IT SIMPLY, HPC IS THE IT OF TOMORROW. »

Born just a few short months ago, the HPC Academy, a nonprofit organization, is distinguished by its commitment to promote the power of HPC technologies to audiences as diverse as academics and small and medium enterprises. An ambitious project, brought about by enthusiasts like Christophe Rodrigues. The Executive Director of the HPC Academy has agreed to answer our questions to give us an update on the association's development and projects.

The creation of the HPC Academy was a surprise for everyone. How have you been received by the HPC community?

It's true that we never communicated publicly on the proposed creation of the association - only a small circle was aware of the approach. We took this initiative to heart and its implementation was a priority. The beginning was difficult. Some wrongly compared our objectives and missions to Teratec, and I will not deny that we had some pressure to make us stop. However, from the beginning, we had the support of many industry leaders, academics and scientifics who were convinced of the merits of this initiative.

How do you explain why this was so difficult?

Both in France and Europe - because the objectives of the HPC Academy are European - the HPC community is a small one where everyone knows each other. Therefore, it's easy to circulate false information in order to do harm, so that the objectives of the HPC Academy may not have been fully understood by everyone. But that has not dampened our motivation. On the contrary, we are totally committed to carry out the missions that we have set. The Academy is now developping fast enough that there's now nothing to envy of other initiatives.

Don't you think that this distrust of the association is because few people knew you? Tell us more about you...

It's possible! However the president of the association, Frédéric Milliot has already had extensive experience and recognition in the community. For my part, I spent more than ten years analyzing IT solutions for high performance computing, especially for the press. This experience has given me a good knowledge of the market. I then headed to an editorial communication agency for several years before joining Digitechnic four years ago. Finally, today our board has eminent personalities from the community that give credibility to our ambitions and our relevance.

«WE IMPLEMENT OPERATIONS WITH STUDENTS FROM PARTNER SCHOOLS. THE ESIEA WILL BE THE FIRST TO BENEFIT FROM THIS INITIATIVE. »

As to that, how the association is it structured?

The HPC Academy has an executive committee to implement the operational functioning of the association. We have a very important feature: no industry research center, public or government agency members sit on the executive committee. The HPC Academy is thus entirely independent. This operational governance is complemented by three specialized advisory committees to direct our missions.

First, the Education and Research Committee aims to define and implement projects to the academic sphere. This council is composed of six individuals recognized in their fields.

A second committee, in charge of the promotion of HPC, includes all of our member partners: manufacturers, publishers, members of the press or academic. This task force, which is very active in the association, is consulted throughout the year on the activities and projects to implement.

Finally, a Scientific Committee, whose composition will be announced in a few weeks, will have the responsibility of initiating collaborative research projects whose results will then be donated to the open source community.

What concrete actions have been undertaken by the association in Europe?

Our primary mission is to introduce the academic community (teachers and students) to supercomputing and all its applications - far beyond the simple simulation - our actions are primarily oriented towards this goal. For example, we implement operations with students from partner schools. The ESIEA will be the first to benefit from this initiative. One Thursday per month, the HPC Academy teams will host a day about HPC technologies. A speaker (industrial, scientific, academic ...) will introduce students to a technology, a solution, a use case or a parallel programming language. Entitled, «HPC University», this program is free and can be replicated in any school of higher education in Europe.

We also put the finishing touches to a research and collaborative development program called «HPC Camp ‘16.» It consists of ten engineering students interning for six months and another twenty working remotely. Their mission will be to develop an application project and to offer the results to the community. For this we have identified five themes for the future: Machine Learning Deep Learning, HPC Cloud, Cryptology and Security, Financial calculation. On this basis, our Board of Education and Research should soon approve the selected projects. The purpose of HPC Camp ‘16 is to stimulate innovation among the student community, we hope it will lead some participants to engage in entrepreneurship.

Finally we are preparing many events. The first will take place in the month of January 2016 in Paris and is expected to host nearly 500 students. The theme will be «HPC for Security.» This day will be a day of technical conferences, industrial and scientific, technical workshops with practical implementation of the knowledge acquired during the day. It will end with a nighttime hacking competition (hackathon) on the theme: Security.

«DON'T FORGET THAT THESE STUDENTS ARE THE FUTURE PLAYERS OF THE HPC MARKET.»

The winners will be rewarded. The objective of this initiative is to enable participants to have a day in the best synthesis of technical knowledge and best practices so they can quickly integrate into operational businesses. Don't forget that these students are the future players of the HPC market...

You also offer continuing education courses for companies. Isn't that contrary to the «non-profit» nature of the association?

We do indeed offer a program called «HPC Professional Training». This is intended to enable small businesses to acquire skills in parallel programming and high performance infrastructure. We also extend this education to job seekers during retraining. Our objective is to provide new opportunities for businesses development related directly or indirectly to HPC technologies. We have developed five modules of ten days with different levels of technical requirements, and other courses also will be launched within a few months. It is important to remember that the HPC Academy does not intend to sell these kinds of trainings. So we signed an agreement with the ESIEA School of Engineers in Laval and Paris. It is this school which sells and accredits the curriculum by the HPC Academy. The instructors are recruited and provided by the association. We hope to develop such training in other countries in Europe very quickly.

Conversely, what about the actions towards the uninitiated and even the general public?

Excellent question! We believe that HPC is now out of its ivory tower and that, to put it simply, HPC is the IT of tomorrow. Therefore, we are creating a completely free training we will probably call «Beginning in HPC.» This will be a one to three day conference where we inform both professionals (companies, developers, resellers, etc.) and independents that wish to better understand supercomputing, areas of applications, technologies and existing solutions. This new ambitious program will be offered in January throughout Europe in collaboration with our partner members. It will be organized quarterly with special additional sessions at international events such as ISC, for example.

Last, but not the least, question : How do you finance these projects?

Even for a non-profit organization, the money is crucial. The HPC Academy funding sources are currently limited to professional dues (membership is free for students, teachers and scientists) and donations. Even there, it's a struggle every moment to convince donors to support our initiatives. We don't benefit from any public funds or European subsidies, and the bad publicity made by some did not help us. Despite this, we were able to count on the support of our first partner members, which do not hesitate to make material donations (tax deductible, as well as monet ary donations) to equip our classrooms and our stakeholders. Some have already understood: generosity here doesn't have just one goal, the goal is to contribute to the development of HPC in all its forms, for the good of the community, the economy and European innovation.

books

PROGRAMMING MASSIVELY PARALLEL PROCESSORS: A HANDS-ON APPROACH

David B. Kirk and Wen-Mei W. Hwu

Morgan Kaufmann Publishing

512 pages, 57,73€

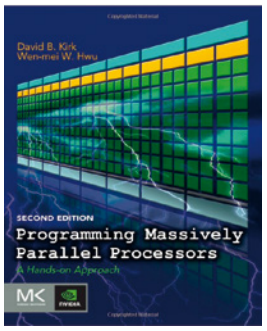

This best-selling guide to CUDA and GPU parallel programming has been revised with more parallel programming examples, commonly-used libraries, and explanations of the latest tools. With these improvements, the book retains its concise, intuitive, practical approach based on years of road-testing in the authors' own parallel computing courses. «Programming Massively Parallel Processors: A Hands-on Approach» shows both student and professional alike the basic concepts of parallel programming and GPU architecture. Various techniques for constructing parallel programs are explored in detail. Case studies demonstrate the development process, which begins with computational thinking and ends with effective and efficient parallel programs. Updates in this edition include: new coverage of CUDA 4.0, improved performance, enhanced development tools, increased hardware support, and more; increased coverage of related technology OpenCL and new material on algorithm patterns, GPU clusters, host programming, and data parallelism; and two new case studies explore the latest applications of CUDA and GPUs for scientific research and high-performance computing.

SHARED MEMORY APPLICATION PROGRAMMING

Victor Alessandrini

Morgan Kaufman Publishing

556 pages, 50,95€

Shared Memory Application Programming presents the key concepts and applications of parallel programming, in an accessible and engaging style applicable to developers across many domains. Multithreaded programming is today a core technology, at the basis of all software develop-

ment projects in any branch of applied computer science. This book guides readers to develop insights about threaded programming and introduces two popular platforms for multicore development: OpenMP and Intel Threading Building Blocks (TBB). Author Victor Alessandrini leverages his rich experience to explain each platform's design strategies, analyzing the focus and strengths underlying their often complementary capabilities, as well as their interoperability. The book develops the essential concepts of thread management and synchronization, and an in-depth discussion of TBB and OpenMP including the latest features in OpenMP 4.0 extensions to ensure readers' skills are fully up to date.

HIGH PERFORMANCE PARALLELISM PEARLS: MULTICORE AND MANY-CORE PROGRAMMING APPROACHES

James Reinders and Jim Jeffers

Morgan Kaufmann Publishing

548 pages, 55,75€

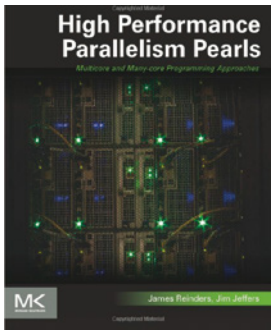

High Performance Parallelism Pearls shows how to leverage parallelism on processors and coprocessors with the same programming - illustrating the most effective ways to better tap the computational potential of systems with Intel Xeon Phi coprocessors and Intel Xeon processors or other multicore processors. The book includes examples of successful programming efforts, drawn from across industries and domains such as chemistry, engineering, and environmental science. Each chapter in this edited work includes detailed explanations of the programming techniques used, while showing high performance results on both Intel Xeon Phi coprocessors and multicore processors. Learn from dozens of new examples and case studies illustrating «success stories» demonstrating not just the features of these powerful systems, but also how to leverage parallelism across these heterogeneous systems.

moocs

PREDICTIVE ANALYTICS

Started on November 25, 2015

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

right competitive strategy to make companies successful, then join us to master the tools of predictive analytics.

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

OPTIMIZING MICROSOFT WINDOWS SERVER STORAGE

Starts on December 15, 2015

Go beyond simple Microsoft Windows Server storage strategies with three key storage optimization features: iSCSI Storage, Storage Spaces, and Data Deduplication. Optimizing Microsoft Windows Server storage is key to keeping pace with your organization's ever increasing data needs. Once you have initially configured your Microsoft Windows Server storage you will want to take advantage of three key features: iSCSI Storage, Storage Spaces and Data Deduplication. iSCSI is a protocol that supports remote access to SCSI-based storage devices over a TCP/IP network. It provides and easy to use alternative to Storage Area Networks (SANs) and can use existing infrastructure. Storage Spaces lets you group physical disks together and present them as a single logical disk. This makes it easy to manage and dynamically allocate storage. Data Deduplication is a service that identifies and removes duplications within data. The goal of Data Deduplication is to maximize the use of disk space. This course follows INF201.21, Implementing Microsoft Windows Server Disks and Volumes but it is not required. Go beyond simple storage strategies and take control of your organization's storage needs! What you'll learn ? How to implement Microsoft Windows Server iSCSI infrastructure to access storage devices on the network. How to implement Microsoft Windows Server Storage Spaces including storage layout, drive allocation and provisioning. How to implement the Microsoft Windows Server Data Deduplication service to maximize the use of disk space.

Length: 6 weeks **Effort:** 2 - 4 hours/week
Price: FREE **Add a Verified Certificate for $49**
Institution: Microsoft **Subject:** Computer Science
Level: Introductory **Languages:** English
Video Transcripts: English
Link : https://www.edx.org/course/optimizing-microsoft-windows-server-microsoft-inf201-22x#!

DATA SCIENCE AND ANALYTICS IN CONTEXT

Learn the foundations of statistical thinking, the power of machine learning, and enabling technologies for data science. Do you want to understand how data science is revolutionizing our world? In this Data Science and Analytics XSeries you will gain a broad understanding of data science fundamentals and learn from real examples how data science can be applied to spark innovation and better decision making. You will find out how to use statistics, machine learning, algorithms, and game-changing technologies, like the Internet of Things to create data-driven solutions and drive improvement in industries ranging from health care, politics, sales, marketing, finance and business. Taught by an interdisciplinary team of Columbia University faculty with acclaimed expertise in data science, thisXSeries is perfect for someone who wants to use cutting-edge data science without the need to learn programming. By taking this Data Science and Analytics XSeries, you will gain critical thinking skills to advance your career, help your organization gain competitive advantage, and understand how data is transforming society today.

Length: 5 weeks per course **Effort:** 7 - 10 hours per week per course **Subject:** Data Analysis & Statistics Computer Science Engineering Business & Management **Institution:** ColumbiaX
Language: English **Video Transcripts: English**
Price: $99 to $149 per course
Link : https://www.edx.org/xseries/data-science-analytics-context

CHIFFRES CLÉS

44 BILLION DOLLARS

Worldwide projected HPC market value by 2020

8,3%

Yearly growth of HPC market

220 BILLION DOLLARS

Compound market value over the 2015-2020 period

Source : Market Research Media

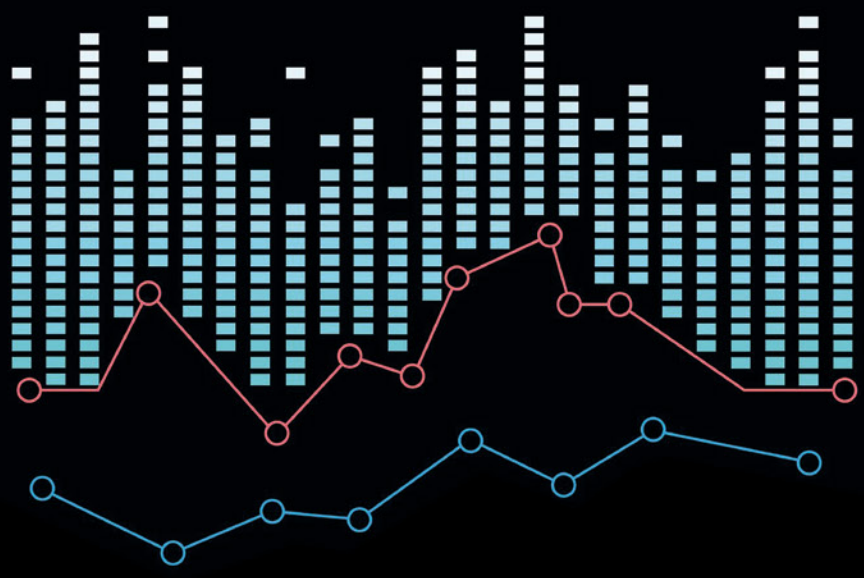

TOP 500
TOP 3

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

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

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

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

GREEN 500
TOP 3

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

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

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

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

HPC REVIEW

SOFTWARE DEFINED DATACENTERS

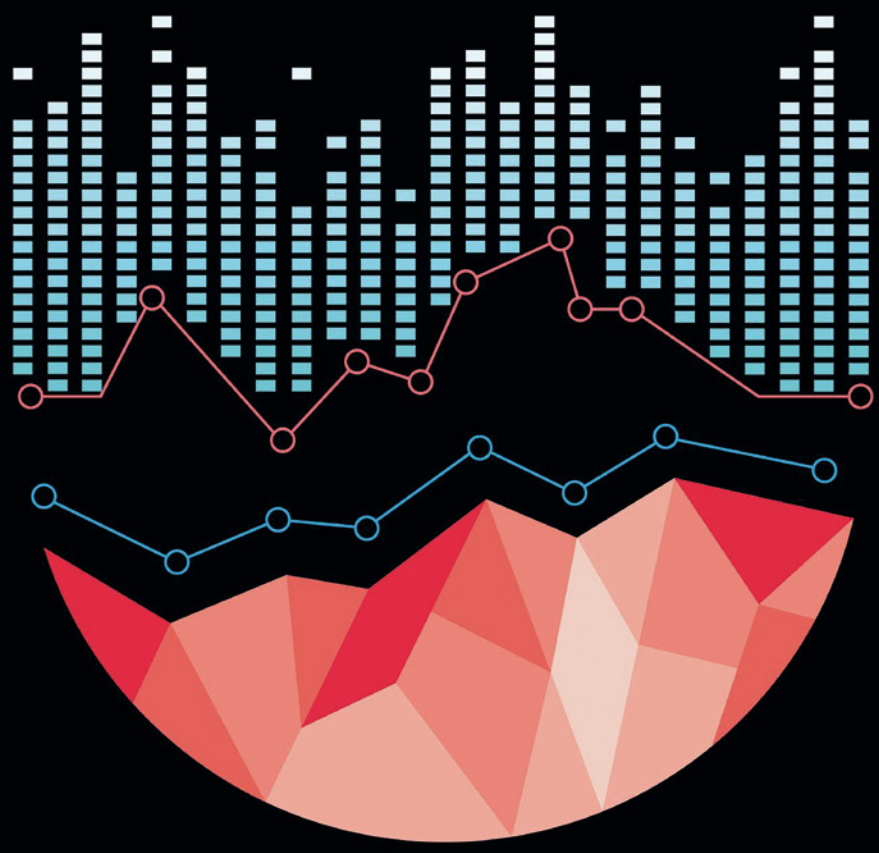

POWERING THE INFRASTRUCTURE OF THE FUTURE

Benefit # 1: **COST** It is more cost-effective to upgrade software than to update hardware infrastructure when development errors occur.

In past issues, HPC Review has covered the unprecedented development and benefits of virtualization of workstations and storage infrastructures in a computing environment, allowing remote supercomputing uses to meet business, academic and research needs. But the changes occur also on the infrastructure side: Datacenters have begun to change dramatically. It is now servers and the network itself to be converted to their virtualized counterparts, with additional key benefits. All industry players offer solutions for linking virtualized environments, with their equivalent data center side, which together to form a single logical link, offering administrators more security, performance and agility. And ultimately, more competitiveness. HPC and end to end virtualization is a winning combination!

NETWORK VIRTUALIZATION: ELIMINATE INCONSISTENCIES

The potential of network virtualization (or SDN) lies in its ability to facilitate the changes in the organization and support different processes where control resources entrusted to business segments which are outside the network or infrastructure enterprise IT. This new operation is complex, but the gains can be substantial. Why try to virtualize a network? To reduce the costs and in order to both simplify and add flexibility which in turns, generates unprecedented usage possibilities. Virtualizing the network is completely synonymous with Software Defined Network allowing to schematize a network on a higher level in the software part by not using physical layers or simple actions. The idea to virtualize the network is not recent. Among the pioneers are Microsoft that in order to reduce networking costs between PCs on a Wan, had the idea to overcome the hardware limits by converting some of the hard-wired tasks into software. Thus the Winmodem software interface was born to replace the dedicated circuit dedicated to communications between PCs on the WAN (usually a modem) which sent the signals via the telephone line. Upgrading the network hardware, usually quite costly, was therefore limited to a software upgrade much more accessible. It alone was a powerful reason enough to switch some of the traditional missions devoted to hardware, to its software defined sibling.

AND THE NETWORK BECAME SOFTWARE

Becoming Software Defined certainly is the future of the network, through which all the digital exchanges take place. Numerous industry actors and network equipment manufacturers of routers, switches, load balancers, firewalls etc, have already jumped on the bandwagon, convinced that tomorrow's network world will be mostly, if not purely software based. One common conception is that the intelligence lies in the software, which in turn can run on almost any hardware with no peculiarities whatsoever – think generic.

Which is precisely what Software Defined X is based upon – abstracting itself from the hardware layer while retaining its former specifics and strong points, and pulling out of the expensive hardware equipment all the organization, orchestration and action automation

Benefit # 2: **SIMPLICITY** - Reduce network complexity through implementation of intelligence in software and speed up implementing new services in the network, simplified deployment by using the software layer.

capabilities, resulting in a network virtualization layer having identical characteristics on which businesses can adapt and reinvent their IT infrastructure. Of course, this transition is of paramount importance for developers and IT administrators so they can operate their infrastructure elements efficiently both on the management side, application side and business side.

In a software defined world, security is of utmost importance. The reason being that when the future's network infrastructure relies on code, any development mistake resulting potentially in a network vulnerability situation can have heavy consequences.

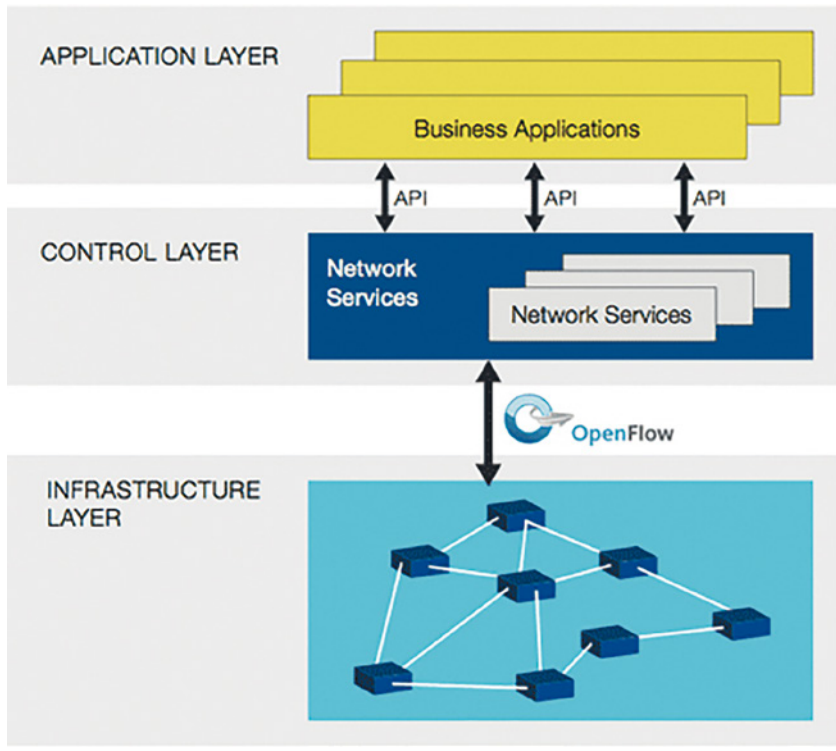

SDN infrastructure proposed by the ONF. There are 3 levels to virtualize networks.

THE SDN VISION : ONE VISION, MULTIPLE APPROACHES

By separating the Data Plane and the Control Plane used respectively for task execution and operation intelligence at the heart of the physical network, three trends have emerged so far.

The control plane controls and centralizes the actions to realize and the operational intelligence. To send commands to the control plane, a virtual layer connect and interact through APIs (Application Programming Interfaces). Which is where the application developers have an important role to play. At this level, an Open Source protocol is also driven by the community : OpenStack, which is used for the data exchanges between the top services layer and the Control Plane.

The requests from these services are translated at the layer control level by simple commands that will be sent to the physical elements scattered anywhere on the network. This communication uses the Open Source Openflow protocol. Dan Pitt, executive director of the OFB (Open Networking Foundation)

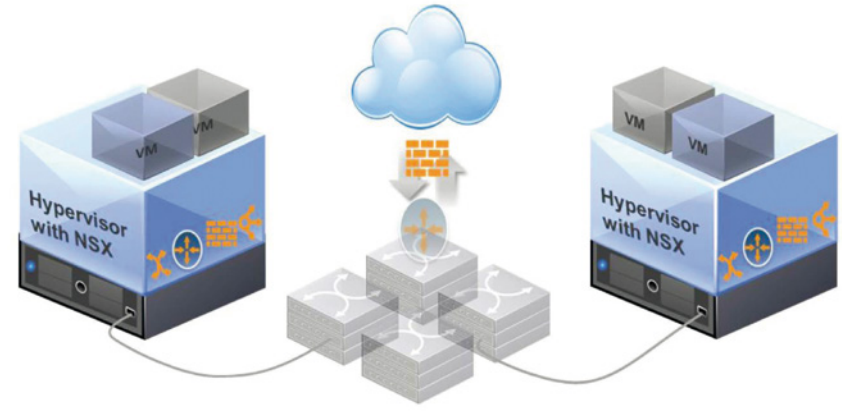

The SDN according VMware at the virtual network NSX include virtual switching level 2 and 3 as well as other features such as firewalls and load balancing.

Benefit # 3: **STANDARD** - arrival of standards will allow companies to mix and match infrastructure elements. It will also allow implementing safety and security measures very early in the development cycle.

Dan Pitt, Executive Director of the UN's engine for the association. He travels the world to evangelize the world of networks and the number of memberships continues to grow.

promotes the use of Open Source protocols around the planet. Beyond this, there are several ways of accomplishing a Software Defined Network software architecture because each network equipment manufacturer has his own vision, a vision strongly influenced by the existing product portfolio.

SAFETY AND SPEED

According to the virtualization pure players, the hypervisor should remain at the center of an SDN architecture. Which in turn, allows to be closer to the virtual machines. Vmware's vision is in fact, to have two hypervisors. One on the server and one on the network, through their dedicated NSX hypervisor. Which in turn, allow to attain both safety and speed, according to VMware.

Besides Open Source protocols and hypervisor based environments, there are other means to achieve the virtualization of a network while keeping the existing network infrastructure. Every network manufacturer is actively working on SDN solutions able to bridge the traditional network infrastructure with an increasingly SDN world. While a full blown SDN network is yet to be achieved, the evolution is going forward, and it is time for companies to begin taking it into account in their infrastructure evolution plans.

SERVER VIRTUALIZATION: CONSOLIDATE AND STREAMLINE

If VMware did not invent virtualization, this company has greatly contributed its dissemination and its modernization. This is a fact: virtualization needs are skyrocketing. Which is not an accident given taking into account the technological developments and the business needs. Therefore the demand is soaring... and it will last! We met Marc Frentzel, Technical Director Europe and Magdeleine Bour-

Benefit # 4: **AGILITY** - The conversion of physical servers into virtual machines provides an unprecedented operational agility that traditional investments do not allow.

goin, Technical Director France to better understand the benefits of server virtualization.

WHY VIRTUALIZE SERVERS?

Virtualization generates many immediately noticeable and quantifiable benefits, since replacement or extension of IT assets are quickly amortized when virtualized in whole or part. It is not the only benefit, since flexibility resource allocation that a virtualized park provides results in a double gain, both operational and financial. According to our interlocutors, the experimentation phase is finally over and organizations can focus further to their digital transformation with improved and robust virtualization infrastructure and solutions. Indeed, the first phase has been crucial to convert the bulk of hardware conversion into virtualized resources. Server virtualization Server is unique since it concerns all sectors of an organization by the use of related technologies on the network (NSX) and storage (vVolumes, vCloud Air).

VIRTUALIZATION IMPROVES THE USE OF IT EQUIPMENT

The second step that companies are preparing to cross is dedicated to the steering and provisioning of application resources. These past are indeed both virtualized third. Their supervision and control required new forms of supervision. This involves understanding the use of monitoring and automation functions to better manage the existing resources in conjunction with a high level of SLA. This technical steering aims to help organizations to keep an eye in the infrastructure economics through the concept of billing that allows to reinvoice resource utilization. On an operational level, let's not forget that 50% of the virtualized workloads in the world are now occupied by SAP, Microsoft and Oracle applications.

OVERPROVISIONING DANGERS

Virtualization has many benefits, but also identified pitfalls. VM deployment has become so simple that administrators often forget how many and which ones they have deployed so far in a giuven point in time. Which results in unused VMs that continue to eat up computing power and resources. The first step of a rearchitected infrastructure is to help regularly control and identify unused

MAGDELEINE BOURGOIN, TECHNICAL DIRECTOR VMWARE FRANCE

«This is not the resource provisioning by itself which raises problems but the risk of oversubscription. It is therefore essential to know the available resources that are not used»

VMWARE LIGHTWAVE: AN OPEN SOURCE INITIATIVE FOR SDDC

Planned to be available within the first half of 2016, VMware launched two initiatives to change and allow data centers developers to design, deploy and administer distributed applications safely. These projects rely on open source evolving technologies, optimized to host and secure applications in virtual container environments à la Docker. These technologies will subsequently be integrated into SDDC solutions to enable companies to deploy containerized applications easier.

The «Lightwave» project is the first management technology for container identity and access. It strengthens security of native cloud applications. Their distributed nature, which may be formed of complex networks of micro services and hundreds or thousands of instances, require that companies ensure supervision identity and access for all components and related users. The project Lightwave adds a layer of additional safety level to the container and beyond its isolation: companies will be able to implement control functions access and management identities over the entire their infrastructure and their application stack, and all stages of the cycle of development.

In addition, the Photon project will allow them to ensure access control so that only approved users are able to use allowed containers on allowed hosts.

Lightwave project functionalities include:

• **Centralized Management Identity** - The Lightwave project functionality offers unique identification, authentication and authorization using names and passwords, from tokens and certificates. Companies will be able to use one solution to secure their native Cloud applications.

• **Supports multitenant architectures** - With Multi-tenant architecture support, the same infrastructure can be used by a variety of applications and teams.

• **Support for Open standards** - The Lightwave project incorporates a multitude of open standards such as Kerberos, LDAP v3, SAML, X.509 and WS-Trust. Which will guarantee interoperability with other standards within the data center.

• **Adapted Scalability** – the world's needs Professional – Project Lightwave is based on a simple and extensible "multimaster» replication model, thus ensuring the put to the required scale while providing high performance.

• **Certification Authority and key management** – The Lightwave will streamline operations requiring certificates and management key on the entire infrastructure.

Benefit # 5: **FLEXIBILITY** - Once the servers and resources are virtualized, their creation and assignment (both internally and externally) according to the business requirements is greatly facilitated. With considerable gains in operating earnings and deadlines.

virtual machines in order to reclaim and re-dispatch the dormant resources on the fly. This ability is crucial in a world where the implementation of VM explodes, and where their Monitoring is not always easy to achieve. To this regard, VMware has a technology (NSX) which adds an arbitration layer to enterprise networks to complete these operations in the most transparent and effective manner.

THE PHOTON PROJECT : A LIGHTWEIGHT LINUX OPERATING SYSTEM AS RUNTIME ENVIRONMENT

The Photon project is a lightweight Linux operating system for container applications, which is a natural complement to the project Lightwave. Optimized for VMware environments vSphere and VMware vCloud Air, the Photon project objectives is to allow businesses to run simultaneously and natively containers and virtual machines on a single platform, and secondly ensure container insulation when executed within virtual machines. This future project will become the perfect blend for container application portability in development and test environments.

The Photon Project features include:

• Supports a variety container formats - The Photon Project ensures Docker containers support, and rkt Garden (Pivotal). Customers and can choose which best fits their needs.

• Heightened Container security - The Photon Project strengthens safety and insulation between container applications within virtual machines. Photon can also control authentication and permissions when integrated with project Lightwave, thus allowing customers to strengthen Security applications down to the layer container itself.

• Flexible scalability and version management - The Photon project provides the administrators and developers extensibility and flexibility to update their runtime host machines. The project relies on the .rpm format of system images version management. It also supports a software package versioning system compatible with YUM, allowing users to ensure granular maintenance.

Benefit # 6: **MONITORING** - The state of the art can today provide a total and complete visibility into the company's use and utilization rates of IT resources.

A MORE FLEXIBLE CONTROL OF RESOURCES

An important concept is that of continuity environments, namely the need control the resources allocated for preserve flexibility without the risk of interrupting exploitation. This ability allows now move workloads without interruption, including those outside the company on partner data centers. What preserves freedom of choice a public cloud and a private cloud. This possibility overflow allows companies keep hold of their costs by transferring in Opex.

«PAY AS YOU GROW «: A RATIONAL INFRASTRUCTURE GROWTH AND COSTS RELATED TO ACTUAL NEEDS

Not so long ago, an organization needed to invest heavily in IT equipment in anticipation of its activity growth. Which induced additional costs related to Installation, skills, dedicated staff and specific expenses (surface occupation, electricity, air conditioning). With the advent of virtualized servers, these costly operations are becoming as trivial as renting a car. With two decisive advantages: on one hand, the lag time between the equipment, operation and amortization stages disappear. A peak of activity, a big project, a large order? Triggering external resources can be done in record time. Does the use of these resources need to be further extended? Easy, simply extend the contract on a weekly or monthly basis. Moreover, it has become a matter of days, if not hours, to provision unprecedented capacity previously only achievable through physical equipment! **SOLANGE BELKHAYAT-FUCHS** & **JOSCELYN FLORES**

HPC REVIEW

THE WORLD'S FIRST ARM BASED HPC CLUSTER

RESEARCH ON COMPUTER ARCHITECTURE BASED ON DISTRIBUTED CLUSTER STRUCTURE

With the fast development of image science and embedded system technique, it is widely accepted that future HPC systems will be limited by their power consumption. The current high performance computing system is a commodity server processors, design for many years to achieve maximum performance, and then it dawned on energy efficiency. In this paper we advocate a different approach: computer architecture based on distributed cluster structure. We introduce the architecture of Tibidabo, the first large-scale HPC cluster built from ARM multicore chips, and a detailed performance and energy efficiency evaluation. We now design experience and to improve the energy efficiency of future HPC systems based on the low power core. The experimental result shows the effectiveness of our methodology, further in-depth research portion is also discussed with case studies.

INTRODUCTION

In High Performance Computing (HPC), there is a continued need for higher computational performance. Science major challenges such as engineering, geophysics, bioinformatics, and other types of compute-intensive applications need a large amount of computing power. On the other hand, energy is becoming one of the most expensive resources, which greatly helps to run a large total cost of super computer facilities. In some cases, the total energy cost over a few years of operation can exceed the cost of the hardware infrastructure acquisition [1]. This trend is not limited to HPC systems, can also be applied to the data centers. Energy efficiency is already a primary concern for the design of any computer system and it is unanimously recognized that reaching the next milestone in supercomputers' performance will be strongly constrained by power. The energy efficiency of a system will define the maximum achievable performance. In this article, we take the first steps in the solution by low power high performance computing system using the embedded and mobile devices. Use the CPU from the domain, however, is a chal-

OUR APPROACH FOR AN HPC SYSTEM IS NOVEL BECAUSE WE ARGUE FOR THE USE OF MOBILE CORES

lenge. Most embedded CPUs lack a vector floating-point unit and their software ecosystem is not tuned for HPC. What makes them particularly interesting is the size and power characteristics which allow for higher packaging density and lower cost. In the following three subsections we further motivate our proposal from several important aspects.

THE ROAD TO EXASCALE AND ARM PROCESSOR

To illustrate our point about the need for low-power processors, let us reverse engineer a theoretical Exaflop supercomputer that has a power budget of 20 MW. An Exaflop machine will require 62.5 million of such cores, independently of how they are packaged together (multicore density, sockets per node). We also assume that only 30-40% of the total power will be actually spent on the cores, the rest going to power supply overhead, interconnect, storage, and memory. That leads to a power budget of 6 MW to 8 MW for 62.5 million cores, which is 0.10 W to 0.13 W per core. Current high performance processors integrating this type of cores require tens of watts at 2 GHz. However, ARM processors, designed for the embedded mobile market, consume less than 0.9 W at that frequency and thus are worth exploring even though they do not yet provide sufficient level of performance and they have a promising roadmap ahead. There is already a significant trend towards using ARM processors in data servers and cloud computing environments [2]. Those workloads are constrained by I/O and memory subsystems, not by CPU performance. Recently, ARM processors are also taking significant steps towards increased double-precision floating-point performance, making them competitive with state-of-the-art server performance. Previous generations of arm application processor has no function Floating-point unit can support HPC1 required throughput and delay. The ARM A9 architecture has an optional VFPv3 floating-point unit [2] and/or neon single instruction multiple data (SIMD) floating point unit [3]. VFPv3 unit is the assembly line, each cycle is able to perform a double add operation, or a MUL (fused multiply accumulation) every two cycles. The neon unit is the SIMD units support only integer and single precision point operand to the HPC itself less attractive. Then, use a double floating point arithmetic instructions per cycle (VFPv3), 1 GHz architecture provide 1-A9 GFLOPS peak. Recently arm architecture (A15 [4] processor has a completely pipelining double-precision floating-point unit and provide 2 GFLOPS 1 GHz per cycle (FMA). The new ARMv8 instruction set, which is being implemented in next-generation ARM cores, namely the Cortex-A50 Series [5], features a 64-bit address space, and adds double-precision to the NEON SIMD ISA, allowing for 4 operations per cycle per unit leading to 4 GFLOPS at 1 GHz.

BELL'S LAW AND CONTRIBUTION

Our approach for an HPC system is novel because we argue for the use of mobile cores. We consider the improvements expected in mobile SoCs in the near future that would make them real candidates for HPC. As Bell's states [6], a new computer class is usually based on lower cost components, which continue to evolve at

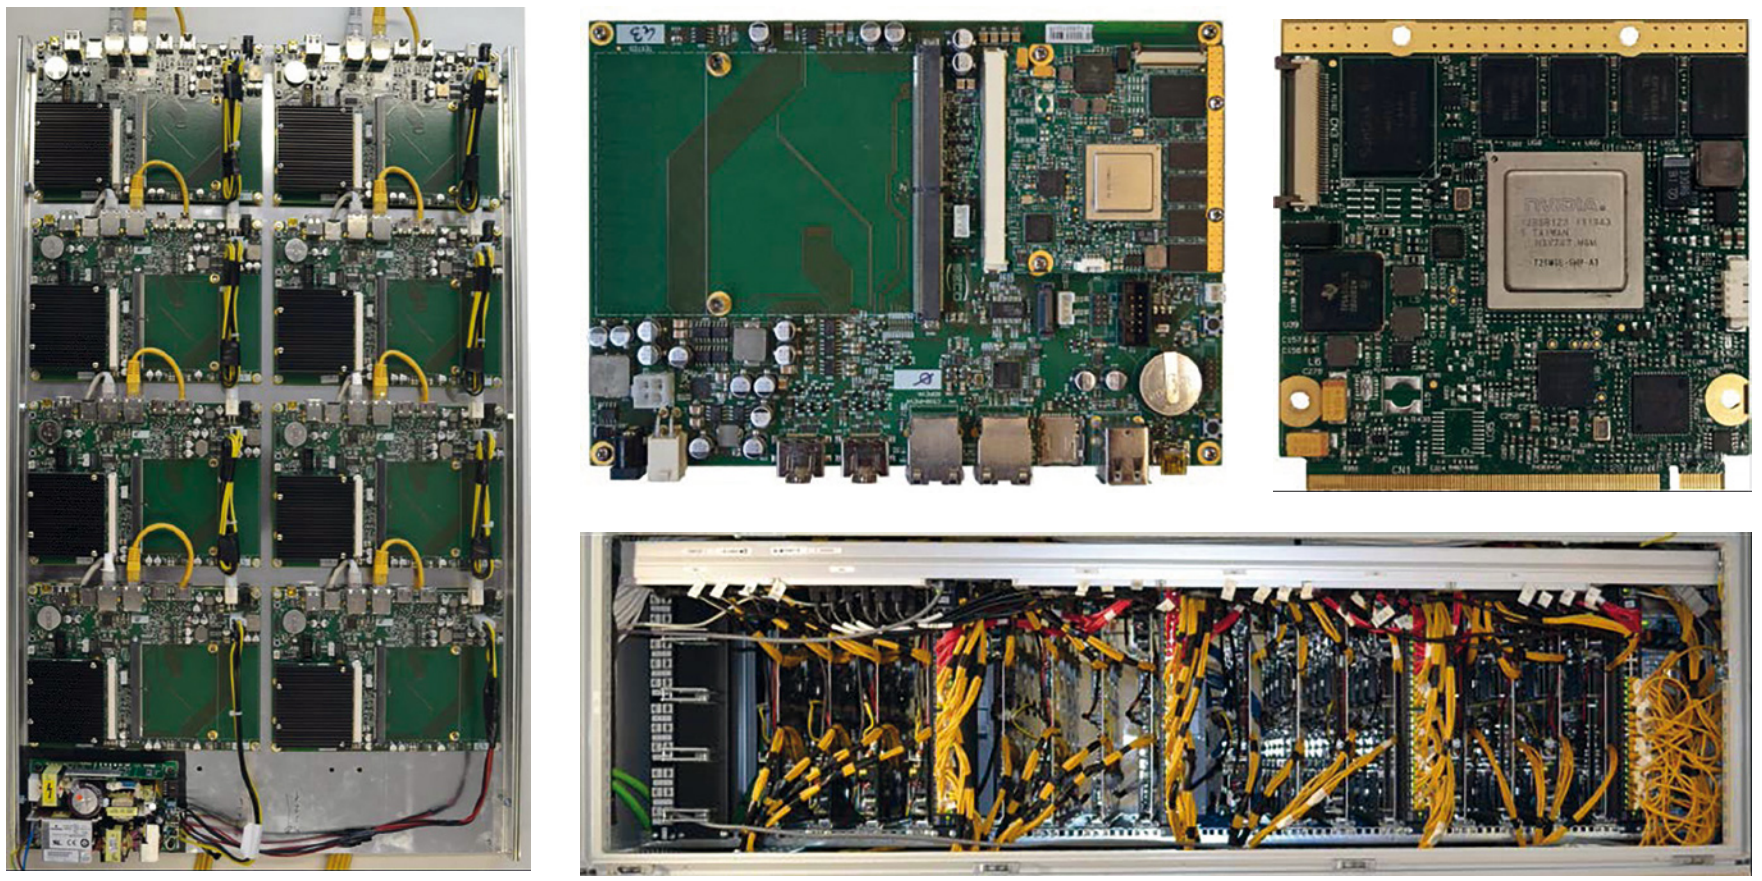

Fig.1: Components of the Tibidabo System

a roughly constant price but with increasing performance from Moore's law. This trend holds today: the class of computing systems on the rise today in HPC is those systems with large numbers of closely-coupled small cores (BlueGene/Q and Xeon Phi systems). From an architectural point of view, we suggest that the in this computing the size of the performance of the class and its growth potential and the evolution of the mobile market. In this paper, we present Tibidabo, an experimental HPC cluster that we built using NVIDIA Tegra2 chips, each featuring a performance-optimized dual-core ARM Cortex-A9 processor. We use the PCIe support in Tegra2 to connect a 1 GbE NIC, and build a tree interconnect with 48-port 1 GbE switches. We do not intend our first prototype to achieve energy efficiency competitive with today's leaders. The purpose of this prototype is to be a proof of concept to demonstrate that building such energy-efficient clusters with mobile processors is possible, and to learn from the experience. On the software side, the goal is to deploy an HPC-ready software stack for ARM-based systems, and to serve as an early application development and tuning vehicle. The contributions of this paper are:

(1) Have design of the first HPC ARM-based cluster architecture, with a complete performance evaluation, energy efficiency evaluation, and comparison with state-of-the-art high-performance architectures.
(2) A power distribution estimation of our ARM cluster. (3) Model-based performance and energy-efficiency projections of a theoretical HPC cluster with a higher multicore density and higher-performance ARM cores. (4) Technology challenges and design guidelines based on our experience to make ARM-based clusters a competitive alternative for HPC.

THE ARM CLUSTER ARCHITECTURE ANALYSIS

The computing chip in the Tibidabo cluster is the Nvidia Tegra2 SoC, with a dual-core ARM Cortex-A9 running at 1 GHz and implemented using TSMC's 40nm LPG performance-optimized process. Tegra2 features a number of application-specific accelerators targeted at the mobile market, such as video and audio encoder/decoder, and image signal processor, but none of these can be used for general-purpose computation and only contribute as a

SoC area overhead. The GPU in Tegra2 does not support general programming models such as CUDA or OpenCL, so it cannot be used for HPC computation either. However, more advanced GPUs actually support these programming models and a variety of HPC systems use them to accelerate certain kind of workloads. Tegra2 is the central part of the Q7 module [7] (See Figure 1(a)). The module also integrates 1 GB of DDR2-667 memory, 16 GB of eMMC storage, a 100 MbE NIC (connected to Tegra2 through USB) and exposes PCIe connectivity to the carrier board. Using Q7 modules allows an easy upgrade when next generation SoCs become available, and reduces the cost of replacement in case of failure. These boards are organized into blades, and each blade hosts 8 nodes and a shared Power Supply Unit (PSU). In total, Tibidabo has 128 nodes and it occupies 42 U standard rack spaces: 32 U for compute blades, 4 U for interconnect switches and 2 U for the file server. These are the basic structure of the proposed system. At the time of writing of this paper this was the only MPI distribution that worked reliably with the SLURM job-manager in our cluster. We use ATLAS 3.9.51 [8] as our linear algebra library. This library is chosen due to the lack of a hand-optimized algebra library for our platform and its ability to auto-tune to the underlying architecture. Applications that need an FFT library rely on FFTW v3.3.1 [9] for the same reasons.

THE EVALUATION AND VALIDATION METHODOLOGY

For single-node energy efficiency, we have measured a single Q7 board and compared the results against a power-optimized Intel Core i7 [10] laptop (Table 1), whose processor chip has a thermal design power of 35 W. Due to the different natures of the laptop and the development board, and in order to give a fair comparison in terms of energy efficiency, we are measuring only the power of components that are necessary for executing the benchmarks, so all unused devices are disabled. On our Q7 board, we disable Ethernet during the benchmarks execution. On the Intel Core i7 platform, graphic output, sound card, touchpad, blue-tooth, WiFi, and all USB devices are disabled, and the corresponding modules are unloaded from the kernel. The hard disk is spun down, and the Ethernet is disabled during the execution of the benchmarks. Multithreading could not be disabled, but all experiments are single-threaded and we set their logical core affinity in all cases. On both platforms benchmarks are compiled with -O3 level of optimization using GCC 4.6.2 compiler.

THE SINGLE NODE PERFORMANCE

We start with the evaluation of the performance and energy efficiency of a single node in our cluster, in order to have a meaningful comparison to other state-of-the-art compute node architectures. In Figure 2 we evaluate the performance of Cortex-A9 floating-point double-precision pipeline using in-house developed micro benchmarks. These benchmarks perform dense double precision floating-point computation with accumulation on arrays of a given size (input parameter) stressing the FPADD and FPMA instructions in a loop. We exploit data reuse by executing the same instruction multiple times on the same elements within one loop iteration. This way we reduce loop condition testing overheads and keep the floating point pipeline as utilized as possible. The purpose is to evaluate if the ARM Cortex-A9 pipeline is capable of achieving the peak performance of 1 FLOP per cycle. Our results show that the Cortex-A9 core achieves the theoretical peak double-precision floating-point performance when the micro benchmark working set in the L1 cache (32 KB).

We also evaluate the effective memory bandwidth using the STREAM benchmark [10]. In this case, the memory bandwidth comparison is not just a core architecture comparison because bandwidth depends mainly on the memory subsystem. However, bandwidth efficiency, which shows the achieved bandwidth out of the theoretical peak, shows to what extent the core, cache hierarchy and on-chip

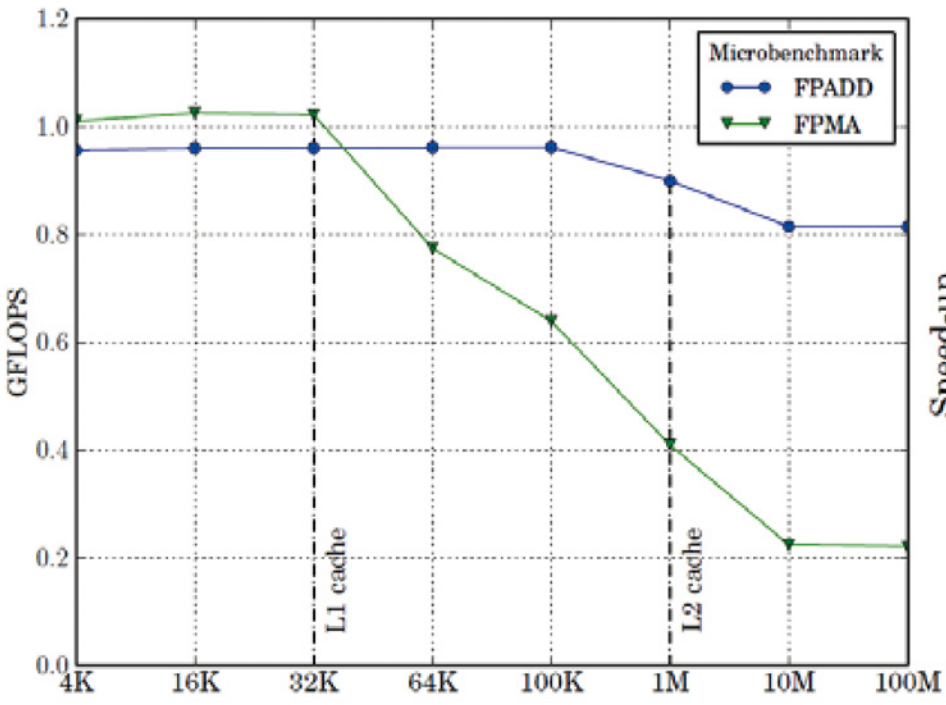

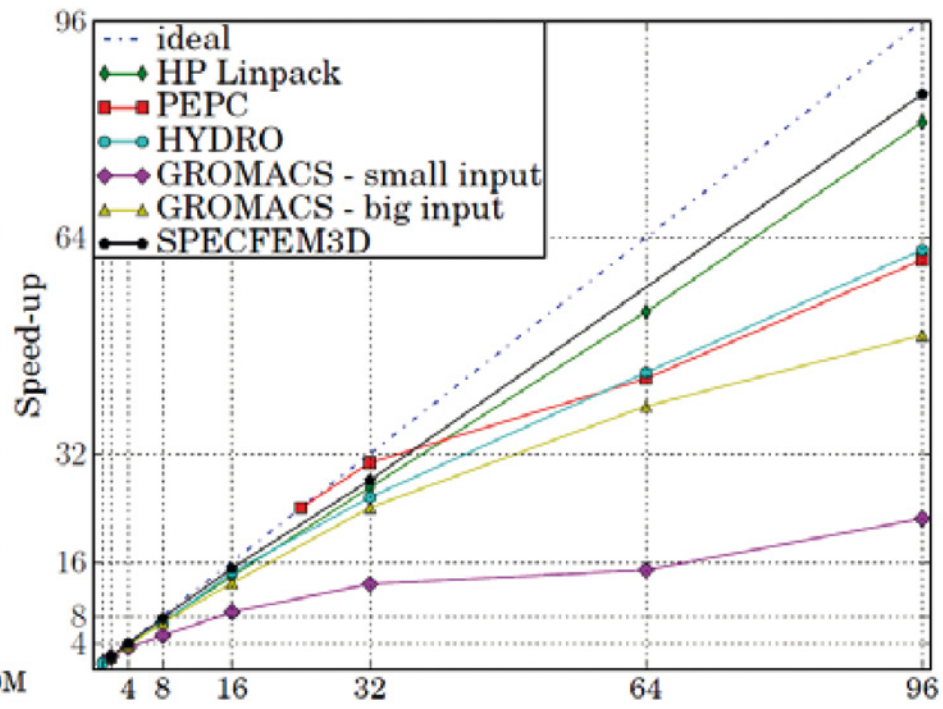

memory controller are able to exploit chip memory bandwidth. We use the largest working set size that in the system. While it is true that the ARM Cortex-A9 platform takes much less power than the Core i7, it also requires a longer runtime, which results in a similar energy consumption the Cortex-A9 platform is between 5% and 18% better. Given that the Core i7 platform is faster, that makes it superior in other metrics such as Energy-Delay. Our single-node performance evaluation shows that the Cortex-A9 is 9 times slower than the Core i7 at their maximum operating frequencies, which means that we need our applications to exploit a minimum of 9 parallel processors in order to achieve competitive time-to-solution. More processing cores in the system mean more need for scalability. In this section we evaluate the performance, energy efficiency and scalability of the whole Tibidabo cluster.

THE CLUSTER ENERGY EFFICIENCY

For both Cortex-A9 and Cortex-A15, the CPU macro power includes the L1 caches, cache coherence unit and L2 controller [11]. Therefore, the increase in power due to a more complex L2 controller and cache coherence unit for a larger multicore are accounted when that power is factored by the number of cores. The memory power is overestimated, so the increased power due to the increased complexity of the memory controller to scale to a higher number of cores is also accounted for the same reason. Furthermore, a Cortex-A9 system cannot address more than 4 GB of memory so, strictly speaking, Cortex-A9 systems with more than 4 GB are not realistic. The remaining power in the compute node is considered to be overhead, and does not change with the number of cores. The board overhead is part of the power of a single node, to which we add the power of the cores, L2 cache and memory. However, we include configurations for higher core counts per chip to show what would be the performance and energy efficiency if Cortex-A9 included large physical address extensions as the Cortex-A15 does to address up to 1 TB of memory. The power model is summarized in these equations:

$$P_{pred} = \frac{n_{tc}}{n_{cpc}} \times \left(\frac{P_{over}}{n_{nin}} + P_{eth} + n_{cpc} \times \left(P_{A9G} + \frac{P_{L2S}}{2} \right) \right)$$

$$P_{over} = P_{tot} - n_{nin} \times \left(P_{mem} + 2P_{A9G} + P_{L2S} + P_{eth} \right)$$

There are still no enclosures announced, and no benchmark reports, but we expect a better performance than ARMv7-based enclosures, due to an improved CPU core architecture and three levels of cache hierarchy. The Calxeda ECX-1000 SoC is built for server workloads: it is a quadcore chip with Cortex-A9 cores running

at 1.4 GHz, 4 MB of L2 cache with ECC protection, a 72-bit memory controller with ECC support, five 10 Gb lanes for connecting with other SoCs, support for 1 GbE and 10 GbE, and SATA 2.0 controllers with support for up to five SATA disks. Unlike ARM-based mobile SoCs, ECX-1000 does not have a power overhead in terms of unnecessary onchip resources and, thus, it seems better suited for energy-efficient HPC. However, to the best of our knowledge, there are neither reported numbers for energy-efficiency of HPL running in a cluster environment (only single node executions) nor scientific applications scalability tests for any of the aforementioned enclosures.

CONCLUSION AND SUMMARY

In this paper we presented Tibidabo, the world's first ARM-based HPC cluster, for which we set up an HPC-ready software stack to execute HPC applications widely used in scientific research such as SPECFEM3D and GROMACS. Tibidabo was built using commodity components that are not designed for HPC. Nevertheless, our prototype cluster achieves 120 MFLOPS/W on HPL, competitive with AMD Operton 6128 and Intel Xeon X5660-based systems. We identified a set of inefficiencies of our design given the components target mobile computing. The main inefficiency is that the power taken by the components required to integrate small low-power dual-core processors sets the high energy efficiency of the cores themselves. We perform a set of simulations to project the energy efficiency of our cluster if we could have used chips featuring higher performance ARM cores and integrating a larger number of them together. Based on these projections, a cluster configuration with 16-core Cortex-A15 chips would be competitive with Sandy Bridge-based homogeneous systems and GPU-accelerated heterogeneous systems in the Green500 list. These encouraging industrial roadmaps, together with research initiatives such as the EU-funded Mont-Blanc project, may lead ARM-based platforms to accomplish the recommendations given in this paper in a near future. In the future, more in-depth research will be conducted and simulated.

Yuwen Zheng is a researcher at the Shandong Women's University, Jinan, 250300, China

REFERENCES

[1] Forshaw, Matthew, A. Stephen McGough, and Nigel Thomas. «Energy-efficient Checkpointing in High-throughput Cycle-stealing Distributed Systems.» Electronic Notes in Theoretical Computer Science 310 (2015).

[2] Mulfari, Davide, Antonio Celesti, and Massimo Villari. «A computer system architecture providing a user-friendly man machine interface for accessing assistive technology in cloud computing.» Journal of Systems and Software 100 (2015): 129-138.

[3] Shukla, Surendra Kumar, C. N. S. Murthy, and P. K. Chande. «Parameter Trade-off And Performance Analysis of Multi-core Architecture.» Progress in Systems Engineering. Springer International Publishing, 2015. 403-409.

[4] Amin, Muhammad Bilal, et al. «Profiling-Based Energy-Aware Recommendation System for Cloud Platforms.» Computer Science and its Applications. Springer Berlin Heidelberg, 2015. 851-859.

[5] Bistouni, Fathollah, and Mohsen Jahanshahi. «Pars network: A multistage interconnection network with fault-tolerance capability.» Journal of Parallel and Distributed Computing 75 (2015): 168-183.

[6] Kaddari, Abdelhak, et al. «A model of service–oriented architecture based on medical activities for the interoperability of health and hospital information systems.» International Journal of Medical Engineering and Informatics 7.1 (2015): 80-100.

[7] Kenkre, Poonam Sinai, Anusha Pai, and Louella Colaco. «Real Time Intrusion Detection and Prevention System.» Proceedings of the 3rd International Conference on Frontiers of Intelligent Computing: Theory and Applications (FICTA) 2014. Springer International Publishing, 2015.

[8] You, Simin, Jianting Zhang, and L. Gruenwald. Scalable and Efficient Spatial Data Management on Multi-Core CPU and GPU Clusters: A Preliminary Implementation based on Impala. Technical Report 2015.

[9] Liu, Xiaoming, and Qisheng Zhao. «Cluster Key Scheme Based on Bilinear Pairing for Wireless Sensor Networks.» Proceedings of the 4th International Conference on Computer Engineering and Networks. Springer International Publishing, 2015.

[10] Singh, Balram, Shankar Singh, and Narendra Kumar Agrawal. «Mobile Agent Paradigm: A Tactic for Distributed Environment.»

[11] Mezhuyev, Vitaliy. «Metamodelling Architecture for Modelling Domains with Different Mathematical Structure.» Advanced Computer and Communication Engineering Technology. Springer International Publishing, 2015. 1049-1056.

ACKNOWLEDGEMENTS

The research work was supported by Shandong Provincial Staff Education office No. 2013-324.

THIERRY AUZELLE
DIRECTOR OF THE STORAGE DIVISION,
ENTERPRISE GROUP, HP FRANCE

HPC REVIEW

ALL FLASH OR HYBRID FLASH

WHY FULL FLASH ENTERPRISE STORAGE SOLUTIONS ARE GAINING MOMENTUM

In terms of data storage, flash storage solutions are progressing irresistibly. The arguments are technical as well as economic.

Many notebook users or ultrabooks are now experiencing the real benefits of flash memory storage, or SSD, replacing hard drives for weight relief, rapid startup boot, and overall performance improvement. The specialized retailer FNAC estimates that in 2 years time, 50% of laptops will be full flash. In business, the use of flash memory to replace disks is already well known for its ability to accelerate dramatically, and immediately, slow applications. Yet according to a recent IDC study with 178 companies and 23 IT service companies (*), only 5.8% of companies surveyed have adopted full flash storage enclosures. They did it for specific needs, independent of the question of the volume. And 30% use mixed or hybrid storage arrays, ie partially flash. According to IDC, hybrid configurations will operate for many years to come.

THE MAIN TRENDS

The capacity of flash storage systems will continue to rise: the first flash modules of 4 terabytes (TB) in only 2.5 inches are appearing. Which is six times better than the most recent record. And the price / capacity ratio will continue to improve. Prices will fall further: in 2018, the price of one terabyte of flash memory will be the same than conventional hard disks. The «100% Flash» approach is already proving more economical. Studies show that, as of now, the overall cost of SSDs, including operation and maintenance, is less than that of conventional discs. One explanation is the impact of the techniques of thin provisioning (forecast management capacity) and deduplication, which optimize capacity and help reduce costs.

«WE MUST STOP THINKING THAT THE FLASH IS EXPENSIVE. **WE MUST CHANGE OUR APPROACH** AND DETECT TARGET APPLICATIONS»

THE PRICE OBJECTION

Many companies suggest the net price of SSDs, would remain high. In fact, one needs to calculate the total cost, up to the scale of the company's data center. It should also enhance the resulting benefits for users. IDC noted: «We must stop thinking that the flash is expensive. We must change our approach and detect target applications.» Is it that we can do without flash? Yes, but the anti-SSD arguments have less and less weight. The flash solutions will gradually impose themselves naturally, because of the advantages in terms of density, performance, and operational simplicity. According to IDC, sales of all flash and hybrid systems in the 1st quarter 2015 were up + 81.6% compared to 2014, reaching 403.1 million dollars.

THE CONCRETE BENEFITS

The boot procedure are significantly accelerated. The standard levels of storage ('gold', 'silver', 'bronze'), always tedious to manage, disappear. Immediately we gain flexibility because the flash configurations are used to create classes of easily manageable services. In addition, the provisioning useful capacity is significantly simpler to orchestrate by administrators.

WHAT IS FLASH USEFUL FOR?

It is sometimes objected that some applications are not necessarily benefitting from flash storage. That's true, but only partially so. Some Big data applications, object storage filesystems or very large scale video streams will not always be adapted to the solutions' flash '. But overall, investment in flash solutions can be quickly profitable. It should in particular benefit the critical datacenter parts that uses conventional hard drives that consume more energy, emit more heat, take up more space, and are disturbed by vibration or pollution. Flash does not suffer from these disadvantages. IDC states: «New technologies optimize the space in the storage array (thin provisioning, deduplication) - which diminishes the TCO (Total cost of ownership).»

IDC continues: «Flash storage is 8 times faster than conventional hard drives and is between 25 and 35 times faster in batch workloads (...) Flash brings a breath of fresh air to the servers, releasing the CPU loads induced by a slow response time in storage. It gives a second life to old servers and extends the life of the servers. It increases the level of consolidation with more databases by server instances, more VMs per physical server, etc. This also results in a 50% decrease in license costs.»

THE TAKE-OFF OF FLASH SOLUTIONS IS CONFIRMED IRRESISTIBLY. **BUSINESS MANAGERS AND EXECUTIVES** QUICKLY UNDERSTAND THE PRACTICAL BENEFITS.

SSDS NEW ENDURANCE

Another objection comes up frequently, involving SSD modules' endurance, which would be limited. It should be noted first that endurance is measured by the number of rewrites per day or DWPD (Data write per day). A recent White Paper (*) compares the different technologies in terms of endurance (write rates on cells or cycles P / E): Enterprise NAND SLC (Single-Level cells) versus MLC Enterprise NAND (Multi-level cells). It shows that two waves of technology have followed:

WAVE 1: this was the emergence of technologies SLC (Single-level cells), in small capacities (200-400 GB). On storage arrays, this test was important because DWPD arrays, such as servers, would manage overall endurance by driving an even distribution of wear cells. This avoided deviate too early still usable cells and thus preserve capacity.

WAVE 2 is the advent of second generation MLC NAND 2 on Multiple-level cell technologies. Here, each component SSD itself manages, independently, the cell wear distribution. As in the case, for example, of the Sandisk FlashGuard technology wich improves the endurance of MLC modules 10 times. This is also true with the 3PAR (HP) high capacity (1.92 TB and 3.84 TB) flash modules. With this enhancement, the cost of storage per TB is now equal to that of standard hard drives.

In conclusion, the take-off of flash solutions is confirmed irresistibly. Business managers and executives quickly understand the practical benefits. The manufacturing costs of NAND 2D / 3D will continue to decline. Traditional HDD drives will quickly be surpassed in terms of ratio performance / capacity. All companies that have adopted the 'full flash' say they are satisfied or very satisfied, and they took this option with a view to 10 years.

(*) White Paper: SSD Endurance. Source: HP (05/2015) http://h20195.www2.hp.com/v2/getpdf.aspx/4AA5-7601ENW.pdf?ver=1.0

NEXT ISSUE

HPC REVIEW

#10
DECEMBER 30th, 2015

VIEWPOINT

INSIDE INTEL'S KNIGHTS LANDING PROCESSOR ARCHITECTURE

HOW TO

IBM WATSON, THE MACHINE CAPABLE OF SOLVING FUTURE PROBLEMS

COVER STORY

INTERNET OF THINGS
BIRTH OF AN INDUSTRIAL ECOSYSTEM

TECH ZONE

2020 : THE RACE TOWARDS EXASCALE

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
subscriptions@hpcreview.com
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
Solange Belkhayat-Fuchs
Marc Paris

ART DIRECTOR
Bertrand Grousset

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