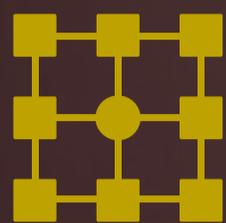


NEWSFEED
IDC HPC SERVER
MARKET
OVERVIEW

LAB REVIEW
LENOVO
THINKPAD
W550S

HOW TO
BUILD
A 1PFLOPS
SUPERCOMPUTER

VIEWPOINT
HOW LUSTRE
ADDS VALUE
TO HPC



HPC REVIEW

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

www.hpcreview.com

HYPERCONVERGED SERVERS

COMPUTE + NETWORKING
+ STORAGE APPLIANCES

BENEFITS AND REALITY CHECK



N°7 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.



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

SHALL HYPERCONVERGENCE RULE ?

OUR COVER STORY THIS MONTH CONCERNS A FUNDAMENTAL ASPECT OF BUSINESS AND DATA CENTER INFRASTRUCTURE. Following a period during which convergence has made significant strides and the software defined X approach has made its mark, the industry is slowly but steadily turning to hyperconverged servers, and its many benefits: performance, simplified administration, consistent solutions, reduced energy consumption ... do these benefits justify breaking the bank and throwing away all or part of existing infrastructures?

NOT NECESSARILY, ACCORDING TO SOME INDUSTRY INSIDERS who advocate a reasoned approach based on progressive renewal phases. At the end of the day such a significant shift is not only a purely technological choice, but also a means to conduct better and smarter business. In an increasingly fast-paced environment, business productivity has become critical.... so it's only logical that the underlying infrastructure remains consistent with strategic needs and ambitions in a competitive market. Don't hesitate to let us know your thoughts at editorial@hpcreview.com. Happy reading !

HAPPY READING!

HYPERCONVERGED SERVERS



**COMPUTE + NETWORKING
+ STORAGE APPLIANCES**
BENEFITS AND REALITY CHECK

NEWSFEED

**IDC VIEW
ON WORLDWIDE HPC
SERVER**



**BARACK OBAMA,
NCSI INITIATIVE
NEW OCZ NVME SSDS**

LAB REVIEW



THINKPAD W550S

HOW TO

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

VIEWPOINT

**HOW LUSTRE
ADDS VALUE
FOR COMMERCIAL HPC**

TECH ZONE

SSDS FOR BIG DATA

EDITOR'S NOTE
OPINION
SOURCES
THE HPC OBSERVATORY

ASUS[®]
IN SEARCH OF INCREDIBLE



ESC4000 G3

Green500 Champion Successor

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

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

WHEN MASSIVE DATA NEVER BECOMES BIG DATA

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

The PRACE Days conference in Barcelona last year provided powerful reminders that massive data doesn't always become big data — mainly because moving and storing massive data can cost massive money. PRACE is the Partnership for Advanced Computing in Europe, and the 2014 conference was the first to bring together scientific and industrial users of PRACE supercomputers located in major European nations.

Tom Lange of Procter & Gamble, which has used high performance computer (HPC) systems to help design consumer products including Pampers diapers and Pringles potato chips, said that although P&G manufactures billions of Pampers a year and they all have data sensors affixed to them, all this data is deleted once the diapers exit the manufacturing line. «P&G doesn't have any big data,» he stated, explaining that «businesses typically won't pay for storing data that doesn't make

them money.» Someone has to make the economic argument for storing data, especially when it comes in petabyte-size chunks.

This revelation sparked a lively discussion on who gets to decide which data gets saved when the future value of the data isn't known. For businesses like P&G, the decision process is clear: if no one steps forward to pay for storage, the data probably won't be saved. For massive data collected by governments, decisions can be more serendipitous and comprehensive policies are often lacking today.

Because the HPC community confronts very massive data, much of it meeting the «four Vs» definitional criteria for big data (volume, variety, velocity, and value), leading HPC users are already wrestling with issues that may not affect mainstream IT markets for a while. In addition:

— **THE LARGE HADRON COLLIDER AT CERN** generates 1 PB of data per second when it's operating and the upcoming Square Kilome-

STORAGE IS NOT THE ONLY DAUNTING EXPENSE ASSOCIATED WITH MASSIVE DATA. MOVING DATA CAN ALSO BE VERY COSTLY. TECHNICAL EXPERTS SAY THAT A SINGLE COMPUTATION TYPICALLY REQUIRES 1 PICOJoule OF ENERGY, BUT MOVING THE RESULTS OF THE COMPUTATION MAY COST AS MUCH AS 100 PICOJouLES.

ter Array telescope will disgorge 1 EB of data (1,000 PB) per day. In both cases, only a small fraction of the data will be stored. Storing it all would break the budget.

— **CHINA'S INTERNET-OF-THINGS INITIATIVE** is targeting a 10,000-fold data reduction to avoid having to confront 100ZB of data from home sensors alone by 2030. A major strategy of the effort, spearheaded by the Chinese Academy of Sciences, is to create a single megasensor that will do the work of the 40-50 device sensors that would otherwise be expected to inhabit the average Chinese home at that time.

The HPC community has already figured out how to exploit the benefits of Hadoop while evading Hadoop's requirement to store three copies of data. The workaround is to decouple Hadoop from the Hadoop Distributed File System (HDFS) and attach it instead to a fully POSIX-compliant parallel file system.

Storage is not the only daunting expense associated with massive data. Moving data can also be very costly. Technical experts say

that a single computation typically requires 1 picojoule of energy, but moving the results of the computation may cost as much as 100 picojoules. At an average \$1 million per megawatt, skyrocketing energy costs have become the number two issue in the worldwide HPC community, right after the perennial desire for bigger budgets. Hence, curtailing data movement is a bulls-eye target for cost-conscious HPC users in government, academia, and industry.

The HPC community was arguably the original home of big data and still operates at the leading edge of many big data developments, but costs for data movement and storage will continue to act as a brake on the growth of HPC big data. Even with this partial brake, IDC forecasts that the server market for high performance data analysis (big data using HPC) will grow rapidly (23.5% CAGR) to reach \$2.7 billion in 2018 and the related storage market will expand to about \$1.6 billion in the same year.

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

Meet the all-new **Saber 1000**

SATA 3.0 Enterprise SSD series.

Features

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



Delivers great performance for today's typical datacenter applications



Read Cache
& Indexing



Virtual Desktop
Infrastructure
(VDI)



Front-End
Web Server



Media
Streaming



Video on Demand
(VoD)



Decision Support
System (DSS)



Cloud
Infrastructure



Customer
Relationship
Management (CRM)



Online
Archiving



Video Editing/
Photo Sharing



Enterprise Content
Management
(ECM)



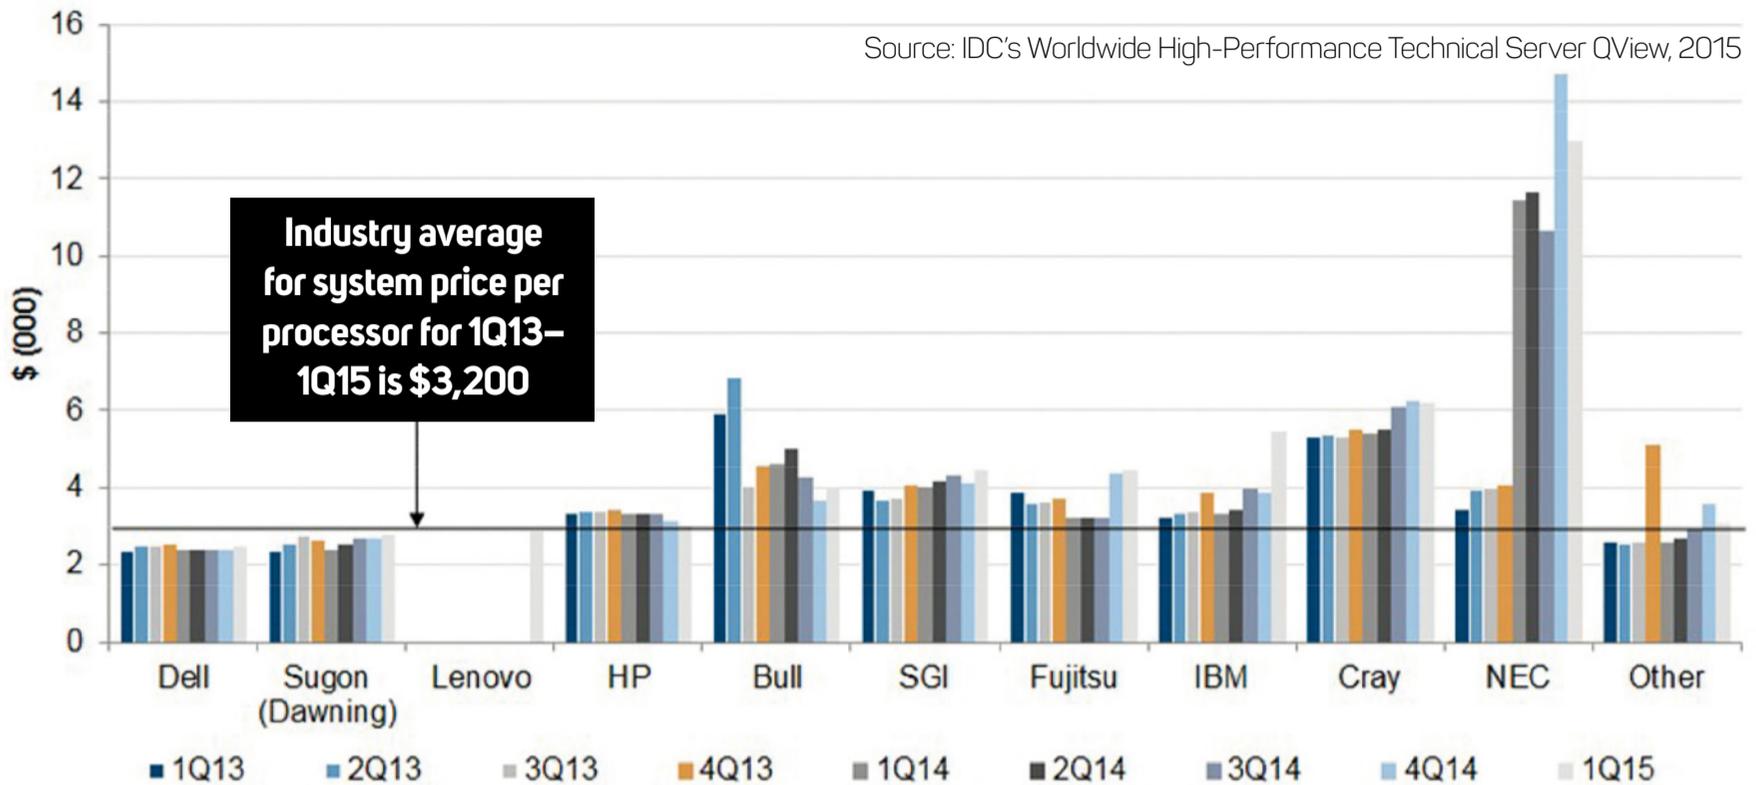
Virtual Tape
Library (VTL)



IDC VIEW

ON WORLDWIDE HPC SERVERS





DC's High Performance Technical Computing group has released its Worldwide HighPerformance Technical Server QView with technical and market data up to first quarter of 2015. Overall, the technical server market grew by more than 10% on a year-over-year basis to over \$2.5 billion for 1Q15.

IDC ANALYSTS NOTE IT IS WORTH SHARING THE FOLLOWING INSIGHTS

The IBM/Lenovo spilt is displaying some degree of conservation of mass. For the first quarter of 2015, the combined technical server revenue of IBM and Lenovo (the combination of IBM Power-based servers and Lenovo's newly acquired x86 servers) was only about 6% less than the total IBM server revenue (Power plus x86) for the first quarter of 2014. At least in this initial phase, Lenovo has done a good job of holding onto much of the transferred IBM x86 server business.

Market share leadership has greatly shifted. Lenovo has moved up to the major leagues, while IBM moves down to triple A. Prior to this quarter, Lenovo sales were too low to move it out of ©2015 IDC #256828 2 the «other» category in the QView vendor listing. In the first quarter of 2015, Lenovo had a solid hold on third place in the technical server market, with revenue behind only HP and Dell. For its part, within the past six months, IBM fell from being the world's second-largest technical server supplier to a distant fourth, well behind Lenovo, which in the first quarter of 2015 has more than three times technical server revenue than IBM. The IBM Power business is showing some strong momentum for the future, such as the large CORAL procurement wins and OpenPOWER Foundation gains, but it will take some time to reach full speed.

Chinese suppliers are making inroads. Chinese technical server vendors appear to be making a move against their competitors in the United States and Japan. In addition to Lenovo's explosive growth — essentially a strong reflection on the ability of the firm to retain IBM x86 custo-

mers — another Chinese technical server supplier, Sugon, saw its first-quarter 2015 year-over-year revenue increase by over 60%. Although Sugon is only about one-seventh the size of Lenovo — and is currently making sales primarily in the Chinese market — its overall growth rate is outpacing most competitors. Lenovo has said publicly that it expects to lose some IBM x86 server business worldwide but make up for this with increased business in China. But Lenovo is pushing hard into the European high-performance computing (HPC) market by joining the European Technology Platform for HPC (ETP4HPC) and establishing a European HPC innovation center in Stuttgart, Germany.

The high end is decidedly not for the faint of heart. The highest end of the technical server sector continues to display a significant degree of choppiness, with huge variations in revenue from one quarter to the next. For example, Cray saw its 4Q14–1Q15 revenue decline by more than 80% (4Q14 was typically very strong for Cray) while realizing a 1Q14–1Q15 annual revenue increase of over 30%. Likewise, Atos-Bull of France saw its 1Q14–1Q15 revenue increase by over 40%, while its 4Q14–1Q15 revenue declined by almost the same percentage.

Figure 1 shows the ratio of total technical systems revenue-to-processor counts in those systems for each of the major technical server suppliers across the past nine quarters and reveals that:

- Post the sale of x86 servers to Lenovo, IBM technical server systems' price per processor has increased significantly (from about \$3,800 per processor to well over \$5,400 per processor), reflecting the increased complexity — and associated cost — inherent in IBM's remaining Power-based server lines.
- Not surprisingly, Lenovo x86 servers heretofore marketed by IBM have a much lower average system price per processor than IBM Power servers, and they are much more in line with other technical server suppliers that rely primarily on x86 components and related technology. Lenovo's economies of scale and singular focus

on x86 servers could make the company a much stronger competitor.

- NEC, Cray, and SGI technical servers — which use more nonstandard, in-house, and proprietary hardware (in addition to their use
- Despite concerns that some white-box suppliers — many of which are currently counted in the «other» category — may be driving the server sector into a race to the bottom, Dell continues to offer the lowest overall system price per processor. Indeed, the average system price per processor for the «other» category has generally increased over the past nine quarters, and vendors are currently realizing almost \$500 more per processor for their systems on average than Dell.

FUTURE OUTLOOK

Data contained in the Worldwide High-Performance Technical Server QView and discussed here indicates that there are a number of emerging trends that bear close observation in the near term as they can have a significant impact on both suppliers and buyers within this sector. Trends to watch include:

- IBM's drive to rely exclusively on the company's more powerful but relatively expensive Powerbased systems to drive future sales
- Lenovo's continued retention of IBM customers — and perhaps even expansion into a new base of customers — for the company's x86 servers, particularly within China and Europe and, to a lesser extent, in the U.S. market; likewise with Sugon as it attempts to expand its reach in China and move out into foreign markets, starting in Asia
- The continued growth of smaller, more differentiated technical server suppliers — like NEC, Cray, SGI, Fujitsu, Hitachi, T-Platforms, and Atos-Bull — to develop and compete effectively with relatively higher-cost and higher-performance systems using more proprietary and custom technology
- The growth of white-box suppliers — that saw a 26% first-quarter 2015 year-on-year growth rate — that appear to be targeting the major brand-name technical server vendors with more than just lower prices.

US
GOVERNMENT
LAUNCHES
**A HPC
STRATEGIC
COMPUTING
INITIATIVE**



OVER THE PAST SIX DECADES, U.S. COMPUTING CAPABILITIES HAVE BEEN MAINTAINED THROUGH CONTINUOUS RESEARCH AND THE DEVELOPMENT AND DEPLOYMENT OF NEW COMPUTING SYSTEMS WITH RAPIDLY INCREASING PERFORMANCE ON APPLICATIONS OF MAJOR SIGNIFICANCE TO GOVERNMENT, INDUSTRY, AND ACADEMIA.

In order to maximize the benefits of HPC for economic competitiveness and scientific discovery, the United States Government will create a coordinated Federal strategy in HPC research, development, and deployment. Investment in HPC has contributed substantially to national economic prosperity and rapidly accelerated scientific discovery. Creating and deploying technology at the leading edge is vital to advancing my Administration's priorities and spurring innovation. Accordingly, this

order establishes the National Strategic Computing Initiative (NSCI). The NSCI is a whole-of-government effort designed to create a cohesive, multi-agency strategic vision and Federal investment strategy, executed in collaboration with industry and academia, to maximize the benefits of HPC for the United States.

Over the past six decades, U.S. computing capabilities have been maintained through continuous research and the development and

deployment of new computing systems with rapidly increasing performance on applications of major significance to government, industry, and academia. Maximizing the benefits of HPC in the coming decades will require an effective national response to increasing demands for computing power, emerging technological challenges and opportunities, and growing economic dependency on and competition with other nations. This national response will require a cohesive, strategic effort within the Federal Government and a close collaboration between the public and private sectors.

With this initiative, the US government wants to ensure sustain and enhance its scientific, technological, and economic leadership position in HPC research, development, and deployment through a coordinated Federal strategy guided by four principles:

- **DEPLOY AND APPLY** new HPC technologies broadly for economic competitiveness and scientific discovery.

TO ACHIEVE THE FIVE STRATEGIC OBJECTIVES, THIS ORDER IDENTIFIES LEAD AGENCIES, FOUNDATIONAL RESEARCH AND DEVELOPMENT AGENCIES, AND DEPLOYMENT AGENCIES.

- **FOSTER PUBLIC-PRIVATE COLLABORATION**, relying on the respective strengths of government, industry, and academia to maximize the benefits of HPC.

- **ADOPT A WHOLE-OF-GOVERNMENT APPROACH** that draws upon the strengths of and seeks cooperation among all executive departments and agencies with significant expertise or equities in HPC while also collaborating with industry and academia.

- **DEVELOP A COMPREHENSIVE TECHNICAL AND SCIENTIFIC APPROACH** to transition HPC research on hardware, system software, development tools, and applications efficiently into development and, ultimately, operations.

Executive departments, agencies, and offices (agencies) participating in the NSCI will pursue five strategic objectives :

- **ACCELERATING DELIVERY OF A CAPABLE EXASCALE COMPUTING SYSTEM** that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.

- **INCREASING COHERENCE BETWEEN THE TECHNOLOGY BASE** used for modeling and simulation and that used for data analytic computing.

- **ESTABLISHING, OVER THE NEXT 15 YEARS,** a viable path forward for future HPC systems

even after the limits of current semiconductor technology are reached (the «post- Moore's Law era»).

- **INCREASING THE CAPACITY AND CAPABILITY** of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development.

- **DEVELOPING AN ENDURING PUBLIC-PRIVATE COLLABORATION** to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors.

ROLES AND RESPONSIBILITIES

To achieve the five strategic objectives, this order identifies lead agencies, foundational research and development agencies, and deployment agencies. Lead agencies are charged with developing and delivering the next generation of integrated HPC capability and will engage in mutually supportive research and development in hardware and software, as well as in developing the workforce to support the objectives of the NSCI. Foundational research and development agencies are charged with fundamental scientific discovery work and associated advances in engineering necessary to support the NSCI objectives. Deployment agencies

THE ASSIGNMENT OF THESE RESPONSIBILITIES REFLECT THE HISTORICAL ROLES THAT EACH OF THE LEAD AGENCIES HAVE PLAYED IN PUSHING THE FRONTIERS OF HPC.

will develop mission-based HPC requirements to influence the early stages of the design of new HPC systems and will seek viewpoints from the private sector and academia on target HPC requirements. These groups may expand to include other government entities as HPC-related mission needs emerge.

LEAD AGENCIES

There are three lead agencies for the NSCI: the Department of Energy (DOE), the Department of Defense (DOD), and the National Science Foundation (NSF). The DOE Office of Science and DOE National Nuclear Security Administration will execute a joint program focused on advanced simulation through a capable exascale computing program emphasizing sustained performance on relevant applications and analytic computing to support their missions. NSF will play a central role in scientific discovery advances, the broader HPC ecosystem for scientific discovery, and workforce development. DOD will focus on data analytic computing to support its mission. The assignment of these responsibilities reflects the historical roles that each of the lead agencies have played in pushing the frontiers of HPC, and will keep the Nation on the forefront of this strategically important field. The lead agencies will also work with the foundational research and development agencies and the deployment agencies to support the objectives of the NSCI and address the wide variety of needs across the Federal Government.

FOUNDATIONAL RESEARCH AND DEVELOPMENT AGENCIES

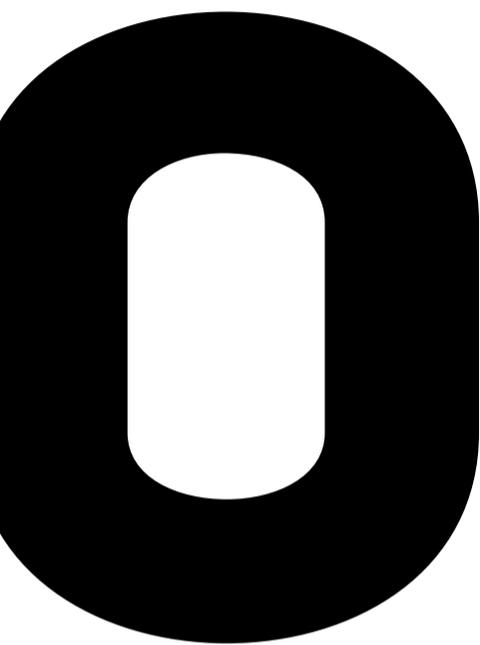
There are two foundational research and development agencies for the NSCI: the Intelligence Advanced Research Projects Activity (IARPA) and the National Institute of Standards and Technology (NIST). IARPA will focus on future computing paradigms offering an alternative to standard semiconductor computing technologies. NIST will focus on measurement science to support future computing technologies. The foundational research and development agencies will coordinate with deployment agencies to enable effective transition of research and development efforts that support the wide variety of requirements across the Federal Government.

DEPLOYMENT AGENCIES

There are five deployment agencies for the NSCI: the National Aeronautics and Space Administration, the Federal Bureau of Investigation, the National Institutes of Health, the Department of Homeland Security, and the National Oceanic and Atmospheric Administration. These agencies may participate in the co-design process to integrate the special requirements of their respective missions and influence the early stages of design of new HPC systems, software, and applications. Agencies will also have the opportunity to participate in testing, supporting workforce development activities, and ensuring effective deployment within their mission contexts.

NEW OCZ NVME SSDS

WIN ENTERPRISE SSD CATEGORY DELIVERING SUPERIOR RANDOM I/O PERFORMANCE, LOW LATENCIES AND ROBUST FEATURES VIA MULTIPLE CONFIGURATIONS



CZ Storage Solutions, a leading provider of high-performance solid-state drives (SSDs) for computing devices and systems, today announced that its new Z-Drive 6000 SSD Series has earned the Best of Computex award in the Enterprise SSD category from industry leading tech publication Tom's Hardware Guide.

This prestigious award recognizing OCZ's innovation and technological advancements was announced June 6th at Computex 2015.

"We are very pleased to win the Best of Computex award in the Enterprise SSD category for our NVMe-compliant Z-Drive 6000 SSD Series," said Ralph Schmitt, CEO for OCZ Storage Solutions. "This award not only represents a validation of our product strategy, but highlights a drive that delivers leading performance and features, and is a compelling solution for our enterprise customers."

CZ's new Z-Drive 6000 SSD Series utilizes NVMe to streamline the storage stack and reduce protocol latency to dramatically boost performance and efficiency over previous Z-Drive generations. PCIe Gen 3 bandwidth allows the Z-Drive 6000 family to achieve leading sustained transfer speeds up to 2.9 GB/s while NVMe efficiency enables the Z-Drive 6000 SSDs to read up to 700,000 I/O requests per second and write up to 160,000 operations per second. The Z-Drive 6000 also delivers consistent and predictable low-latency I/O responses of only 25 microseconds for a 4KB write and 80 microseconds for a 4KB read.

The NVMe-compliant Z-Drive 6000 Series are available in multiple configurations as follows:

- The Z-Drive 6000 SFF Series (read-intensive applications): supports 2.5-inch small form factor (SFF) and usable capacities of 800GB, 1.6TB and 3.2TB.
- The Z-Drive 6300 SFF Series (mixed workload applications): supports 2.5-inch SFF and usable capacities of 800GB, 1.6TB, 3.2TB and 6.4TB.
- The Z-Drive 6300 AIC Series (mixed workload applications): supports Half-Height/Half-Length (HHHL) Add-in Card (AIC) form factors and usable capacities of 800GB, 1.6TB, 3.2TB and 6.4TB.

The Best of Computex Judging Committee was comprised of editors from Tom's Hardware Guide and Tom's IT Pro publications, choosing the most innovative products and components covering graphics technology, motherboards, cooling systems, power supplies, peripherals, storage technology, and general computing. The award showcases vendors pushing the boundaries of technology with a compelling product worthy of the 'Best of Show' recognition.





4 Blades

HighServer XLR4 Blade

Virtualization Server with GPUs / Co-Processors

- High density
- Low power consumption
- Easy management

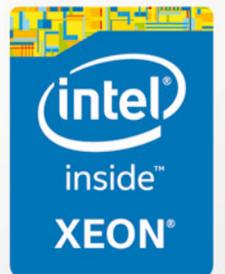
GIGABYTE™

4 GPGPU

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

Available at:

CARRI
Systems



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

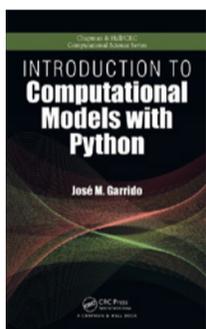


books

INTRODUCTION TO COMPUTATIONAL MODELS WITH PYTHON

Jose M. Garrido

Chapman and Hall/CRC, 466 pages, 49.29£



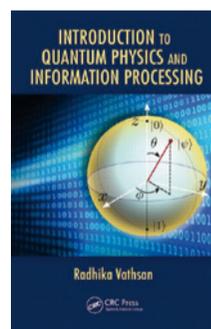
Introduction to Computational Models with Python explains how to implement computational models using the Python programming language. The book's five sections present an overview of problem solving and simple Python programs, introducing the basic models and techniques for designing and implementing problem solutions; programming principles with Python, covering basic programming concepts, data definitions, programming structures with flowcharts and pseudo-code, solving problems, and algorithms; python lists, arrays, basic data structures, object orientation, linked lists, recursion, and running programs under Linux; implementation of computational models with Python using Numpy, with examples and case studies; the modeling of linear optimization problems, from problem

formulation to implementation of computational models. This book provides the foundation for more advanced studies in scientific computing, including parallel computing using MPI, grid computing, and other methods and techniques used in high-performance computing.

INTRODUCTION TO QUANTUM PHYSICS AND INFORMATION PROCESSING

Radhika Vathsan

Chapman and Hall/CRC, 270 pages, 49.29£



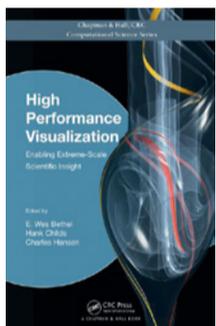
Introduction to Quantum Physics and Information Processing guides beginners in understanding the current state of research in the novel, interdisciplinary area of quantum information. Suitable for undergraduate and beginning graduate students in physics, mathematics, or engineering, the book goes deep into issues of quantum theory without raising the technical level too much. The text begins with the basics of quantum mechanics required to understand how two-level systems are used as



qubits. It goes on to show how quantum properties are exploited in devising algorithms for problems that are more efficient than the classical counterpart. It then explores more sophisticated notions that form the backbone of quantum information theory. Requiring no background in quantum physics, this text prepares readers to follow more advanced books and research material in this rapidly growing field. Examples, detailed discussions, exercises, and problems facilitate a thorough, real-world understanding of quantum information.

HIGH PERFORMANCE VISUALIZATION: ENABLING EXTREME-SCALE SCIENTIFIC INSIGHT

E. Wes Bethel, Hank Childs, Charles Hansen
Chapman and Hall/CRC, 520 pages, 59.49£



Visualization and analysis tools, techniques, and algorithms have undergone a rapid evolution in recent decades to accommodate explosive growth in data size and complexity and to exploit emerging multi- and many-core computational platforms. This book focuses on the subset of scientific visualization concerned with algorithm design, implementation, and optimization for use on today's largest computational platforms. After introducing the fundamental concepts of parallel visualization, the book explores approaches to accelerate visualization and analysis operations on high performance computing platforms. Looking to the future and anticipating changes to computational platforms in the transition from the petascale to exascale regime, it presents the main research challenges and describes several contemporary, high performance visualization implementations. Reflecting major concepts in high performance visualization, this book unifies a large and diverse body of computer science research, development, and practical applications at the intersection of scientific visualization, large data, and high performance computing trends.

mooocs



INTRODUCTION TO C++

Get a brief introduction to the C++ language from the experts at Microsoft. C++ is a general purpose programming language that supports various computer programming models such as object-oriented programming and generic programming. It was created by Bjarne Stroustrup and, "Its main purpose was to make writing good programs easier and more pleasant for the individual programmer.". By learning C++, you can create applications that will run on a wide variety of hardware platforms such as personal computers running Windows, Linux, UNIX, and Mac OS X, as well as small form factor hardware such as IoT devices like the Raspberry PI and Arduino-based boards. What you'll learn : C++ Syntax, C++ Language Fundamentals, How to Create Functions in C++. Prepare yourself for intermediate and advanced C++ topics in follow-up courses taught by Microsoft

Length: 4 weeks **Effort:** 3-5 hours/week

Price: FREE / Add a Verified Certificate for \$90

Institution: Microsoft **Subject:** Computer Science

Level: Introductory **Languages:** English

Video Transcripts: English **Link :** <https://www.edx.org/course/introduction-c-microsoft-dev210x>



ENGINEERING SOFTWARE AS A SERVICE (SAAS), PART 1

UC BerkeleyX

Learn software engineering fundamentals using Agile techniques to develop Software as a Service (SaaS) using Ruby on Rails. This intermediate SaaS course uncovers how to code long-lasting software using highly-productive Agile techniques to develop Software as a Service (SaaS) using Ruby on Rails. Learners will understand the new challenges and opportunities of SaaS versus shrink-wrapped software. They will understand and apply fundamental programming techniques to the design, development, testing, and public cloud deployment of a simple SaaS application. Using best-of-breed tools that support modern development techniques including behavior-driven design, user stories, test-driven development, velocity, and pair programming, learners will see how modern programming language features like

metaprogramming and reflection can improve productivity and code maintainability. Weekly coding projects and quizzes will be part of the learning experience in this SaaS course. Those who successfully complete the assignments and earn a passing grade can get an honor code certificate or verified certificate from BerkeleyX. The videos and homework assignments used in this offering of the course were revised in October 2013. The new class also includes embedded live chat with Teaching Assistants and other students and opportunities for remote pair programming with other students.

Length: 9 weeks **Effort:** 12 hours/week

Price: FREE / Add a Verified Certificate for \$49

Subject: Computer Science **Level:** Intermediate

Languages: English **Video Transcripts:** English

Link: <https://www.edx.org/course/engineering-software-service-saas-part-1-uc-berkeleyx-cs169-1x#!>

Starts on November 2, 2015

PARADIGMS OF COMPUTER PROGRAMMING – ABSTRACTION AND CONCURRENCY

LouvainX

This course covers data abstraction, state, and deterministic dataflow in a unified framework with practical code exercises. Louv1.2x and its predecessor Louv1.1x together give an introduction to all major programming concepts, techniques, and paradigms in a unified framework. We cover the three main programming paradigms: functional, object-oriented, and declarative dataflow. The two courses are targeted toward people with a basic knowledge of programming. It will be most useful to beginning programming students, but the unconventional approach should be insightful even to seasoned professionals. Louv1.1x (Fundamentals) covers functional programming, its techniques and its data structures.

You'll use simple formal semantics for all concepts, and see those concepts illustrated with practical code that runs on the accompanying open-source platform, the Mozart Programming System. Louv1.2x (Abstraction and Concurrency) covers data abstraction, state, and concurrency. You'll learn the four ways to do data abstraction and discuss the trade-offs between objects and abstract data types. You'll be exposed to deterministic dataflow, the most useful paradigm for concurrent programming, and learn how it avoids race conditions.

Length: 6 weeks **Effort:** 4 hours/week

Price: FREE / Add a Verified Certificate for \$100

Subject: Computer Science **Level:** Advanced

Languages: English **Video Transcripts:** English

Link: <https://www.edx.org/course/paradigms-computer-programming-louvainx-louv1-2x-0#!>

**CHIFFRES
CLÉS**

**44 BILLION
DOLLARS**

Worldwide projected
HPC market value by 2020

8,3%

Yearly growth of HPC market

**220 BILLION
DOLLARS**

Compound market value
over the 2015-2020 period

Source : Market Research Media



**TOP 500
TOP 3**

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

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

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

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

**GREEN 500
TOP 3**

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

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

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

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

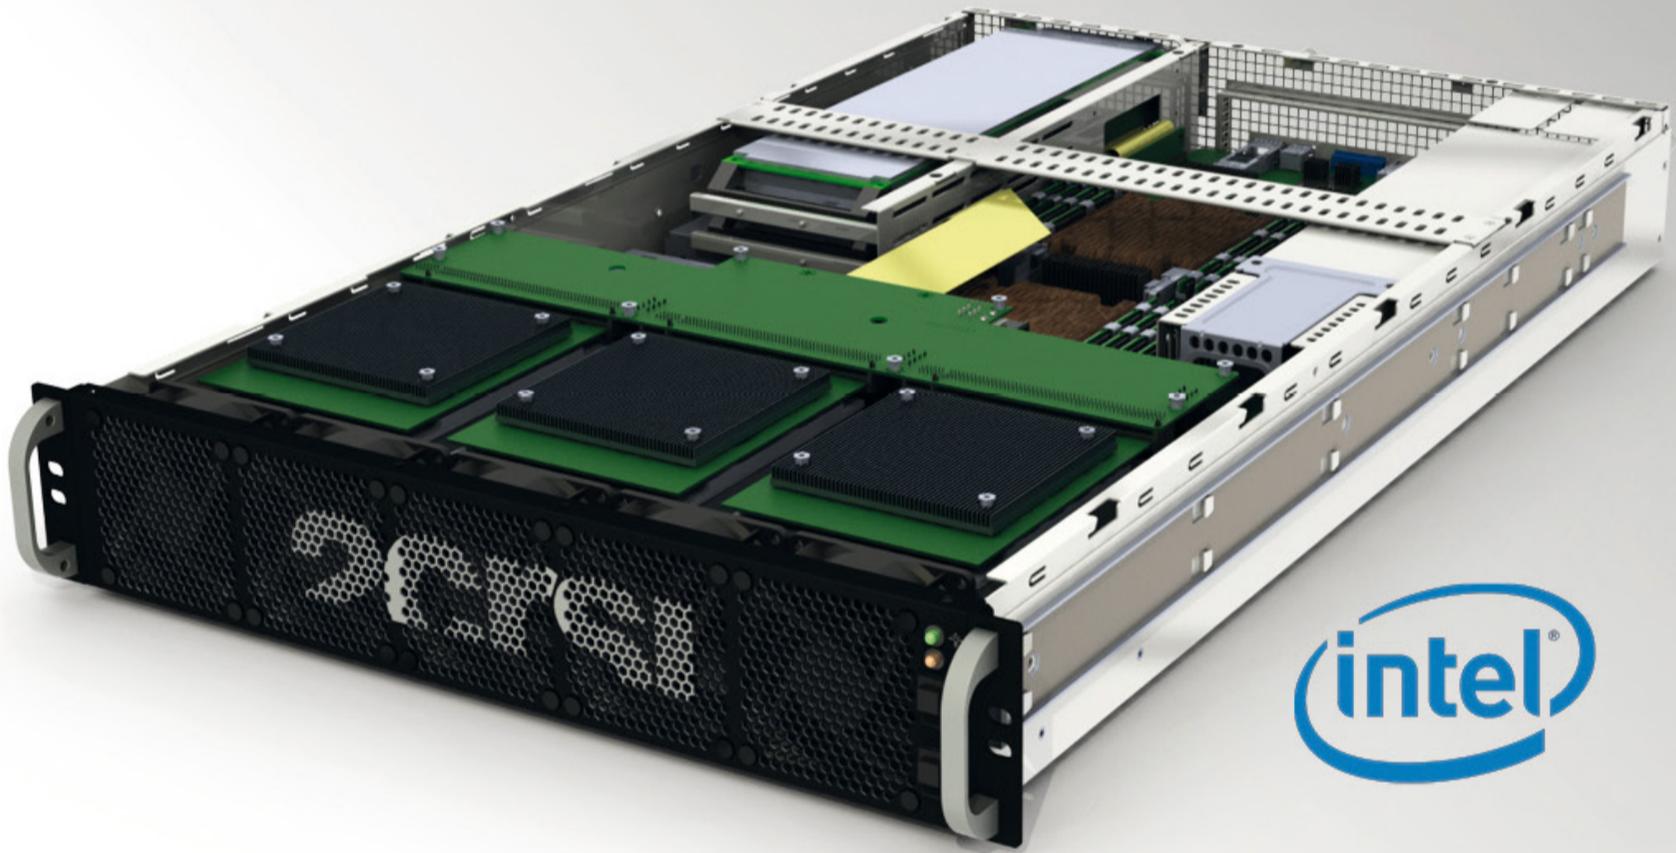
HEXAPHI

“*Innovating in computing...*”



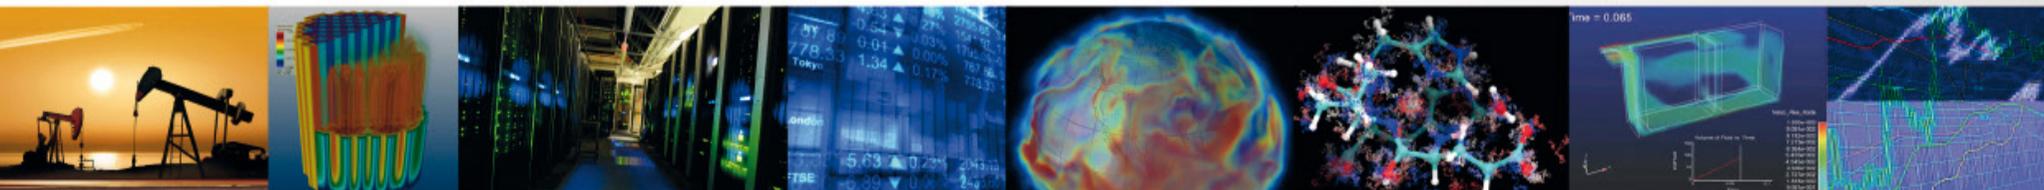
Découvrez la puissance HEXAPHI et sa performance de calcul pouvant atteindre 8 Tflops au sein d'un chassis 2U de 2000W.

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



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

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



HPC@2crsi.com

www.2crsi.com





HYPERCONVERGED SERVERS

COMPUTE + NETWORKING
+ STORAGE APPLIANCES

BENEFITS AND REALITY CHECK





HYPERCONVERGED SERVERS ARE MODULAR SYSTEMS

**DESIGNED TO EVOLVE
OVER TIME DEPENDING
ON LOAD AND NEEDS.
DISCOVER THE
ADVANTAGES AND
DISADVANTAGES OF
THESE SYSTEMS.**



HYPERCONVERGED SYSTEMS ARE A NATURAL EVOLUTION OF THE TRADITIONAL IT INFRASTRUCTURE, WHICH IS USUALLY MADE UP OF SILOS ACCORDING TO BUSINESS AND OPERATIONAL NEEDS.

If you are in the acquisition phase of storage or server platforms, it's impossible not to trip over the recent offers for hyperconverged systems. Your first questions might be: What are the advantages of hyperconverged systems, and how are they different from converged systems? The best questions to ask are: What are converged systems, and how do they differ from traditional infrastructure systems?

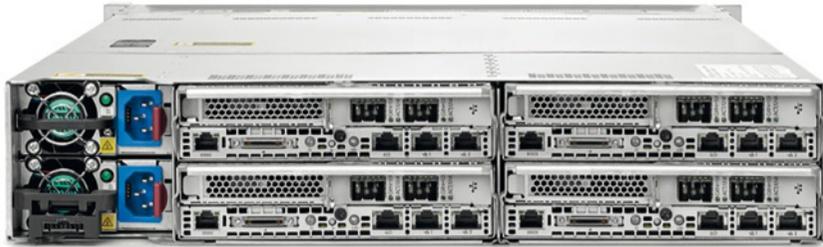
HOW DO HYPERCONVERGED SYSTEMS DIFFER FROM TRADITIONAL INFRASTRUCTURE SYSTEMS?

Hyperconverged systems are a natural evolution of the traditional IT infrastructure, which is usually made up of silos according to business and operational needs. This results in

groups and separate administrative systems for storage, the servers and the network. The group responsible for storage, for example, manages the purchase, supply, and storage support infrastructure, but is not necessarily concerned with network or enterprise server infrastructure issues. The same situation exists for the servers and the network. The relatively recent concept of hyperconverged systems combines two or more of these infrastructure components as a solution.

AT WHAT POINT SHOULD YOU BECOME INTERESTED IN A HYPERCONVERGED SYSTEM?

During the evolution of a business' infrastructure, such as for a data center, there are several pivotal moments that require consideration about its evolution and composition.



According to recognized needs, the identified uses and their alignment with the business' strategy, the path to take will not be the same in all cases. To summarize, there are essentially three options available for a company: expand, consolidate or renew the existing with one or more hyperconverged systems in order to, and this is essential, support and serve the business and its evolution.

FIRST OPTION: EXPAND THE EXISTING

The first option is the expansion of existing infrastructure by adding one or more hyperconverged systems, which means retaining all or part of the existing infrastructure. This approach offers the advantage of not questioning the existing infrastructure and simply adding new network and storage resources. However, keep the consistency of the infrastructure in mind. The risk can be an unbalanced architecture at two distant performance levels from each other, making load balancing tricky. There two aspects to take into consideration: the network of a hyperconverged system is often at 10 Gbit/s, where the majority of firms still operate at 1 Gbit/s. The bottleneck is sudden as the entire network operates at only a tenth of the optimal performance of a new generation appliance. On the storage side it's not much better, although being contained within the same server, hybrid or full flash, can significantly improve the flow rates and response times. Sometimes there are unwelcome surprises due to the heterogeneity and length of storage devices, as demonstrated by the findings of the DataCore survey SDS (see inset). Conversely, the addition of more powerful recent equipment can slow down applications and critical business data, requiring faster access and treatment such as VDI.

ADVANTAGES

- No need to interrupt activity
- Performance gain for storage and applications

DISADVANTAGES

- Risk of imbalanced network flows
- Software Licenses to be considered (VMware)

SECOND OPTION: CONSOLIDATE THE EXISTING INFRASTRUCTURE

When it comes to replacing all or part of a company's servers because of obsolescence or performance reasons, one of the possibilities offered by hyperconverged appliances is allowing their virtualization. The advantages are many: replacing a physical server with a hosted virtual machine, providing a new way of administration, more flexible and unified through a hypervisor. The benefits are also physical(space gained) and economic (reduced energy consumption). Taking into account savings on amortization periods facilitates calculations and comparisons against a hyperconverged server. The process of backup and disaster recovery is thereby also simplified, virtual machines have replaced physical servers by being centralized in a hyperconverged appliance. In this case, it enforces regular backups or replication in a disaster recovery scenario (DRP) or Business Continuity (BCP).

ADVANTAGES

- Centralised administration
- Saving floor space
- Energy savings (performance, cooling)

DISADVANTAGES

- The hyperconverged server represents a single point of failure
- Implementation of emergency backup and/or replication

THIRD OPTION: RENEW THE EXISTING

A hyperconverged approach can selectively replace certain application servers depending on whether they are job or business oriented. It's not uncommon that a server is traditio-



nally dedicated to each application, such as a Microsoft Exchange mail server, for example. P2V Migration (Physical to Virtual) of such a server within a hyperconverged appliance ensures its continuity without requiring reinstallation (the included migration tools abound) while benefiting from the material gain of the appliance. It is therefore possible to gradually replace physical servers by their equivalent in virtual machines. With one caveat, however. If this applies to the vast majority of x86 servers and applications, it does not hold true for «Legacy» servers based on AS/400 hardware, for example.

ADVANTAGES

- Selective P2V Migration spread over time
- Progressive decommission of physical servers
- Intrinsic performance gain

DESADVANTAGES

- Some «Legacy» applications will not tolerate virtualized operation (AS/400, etc.).



HOW HYPERCONVERGED SERVERS DIFFER FROM CONVERGED SYSTEMS

While converged systems are separate components designed to work well together, hyperconverged servers are modular systems designed to evolve by adding additional modules. These systems are designed primarily around storage and computation in a single chassis of x86 servers interconnected by high-speed 10 Gbit Ethernet network interfaces. It's not just a stack of material in a single chassis. The two main differences are 1) the software-defined storage (SDS), which benefits from the flexibility and



performance of the underlying hardware. and 2) the management layer, which controls all the storage resources, calculation and networking from a single point of administration. A third factor relates to scalability, or according to the long-term planning, scale-out of the capacity, that is to say to changing either by adding existing modules, or by interconnecting servers together to form a logical aggregation of multiple hyperconverged servers, they can contain each multiple server blades with their processor, memory, storage, and network interfaces, as offered with its range of HP servers HyperConverged Systems HC-200.

IMPLEMENTATION: EXPRESS!

Another aspect is often overlooked, but it is important: all manufacturers offer hyperconverged deals whose components have been carefully selected and their operation validated. This is not only to propose an offer a single support service, but also to simplify deployment. All the hyperconverged servers are preset and pre-installed to be deployed within fifteen minutes! An IP address for the start, followed by some identifiers is generally enough at a minimal level to be operational on the business level and the current Active Directory . The next step consists of installing or migrating application software and to adjust where necessary, their levels of tolerance/redundancy and charge balancing. With a single dashboard to administer all of it.

TWO WAYS OF UNDERSTANDING HYPERCONVERGENCE

The principle path to getting equipped with a hyperconverged server consists of outfit yourself with a prefabricated system such as those



ON PAPER, HYPERCONVERGED SERVERS HAVE A LOT TO OFFER. THEY ARE NOT, HOWEVER, EXEMPT FROM INCONVENIENCES.



offered by HP, Nutanix, SimpliVity or Scale Computing. It's the simplest and least risky way advised if your activity depends on it. In case you wish to create a custom solution, you can take the server's existing hardware base (available at Supermicro or CARRI Systems) and combine software approaches to it and thus, create your own hyperconverged server. Software offerings such as VMware's VSAN and HP's StoreVirtual are intended for customers wanting more control in the design of their hyperconverged system. With this self-service approach, you will be able to select your server's manufacturer and your preferred software setup, as long as the hardware is supported by the software vendor.

DISADVANTAGES OF HYPERCONVERGED SYSTEMS

On paper, hyperconverged servers have a lot to offer. They are not, however, exempt from inconveniences. Their «black box» aspect,

combined hardware and software prevent bringing substantial modifications to the latter. If the manufacturers are eager to propose the most balanced appliances possible, the possibilities of configurations to gain performance are limited. Similarly, if the appliance has a functional imbalance, in computing power or regarding storage, it's difficult to correct it other than by adding another appliance to compensate for the defect. A default which may not be initially evident, but becomes apparent in the course of operation, once application migrations and virtualization of physical servers have already been made.

CONCLUSION

Despite these reservations, hyperconverged servers have valuable qualities vis-à-vis the modular approach which have prevailed up to now. Among its decisive advantages is the simplification of the infrastructure, the coherence favoring performance and the single administration interface, as well as lower general and administrative expenses and simplified vendor management for highly virtualized environments. Before taking the plunge, it is advised to verify possible gaps in relation to your existing operating environment in order to avoid unpleasant surprises. **JOSCELYN FLORES**



HPCToday

The global media in high-performance IT

Subscribe now > for FREE

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



+ get yourself instant access to exclusive contents and services

www.hpctoday.com



LAB REVIEW

HOW WE TEST

HPC LABS

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

HPCBENCH SOLUTIONS

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

HPC BENCH
GLOBAL INDEX



9 108

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

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



**A TECHNICAL
RECOGNITION
AWARD**



THINKPAD W550S



The Lenovo ThinkPad W550s is a mobile workstation for those who need high-end performance in a package that's easy to carry and travel with. Featuring an Intel Core i7 processor, Nvidia Quadro graphics, and an impressive 3K display, the ThinkPad W550s is a good blend of mobility and performance. The balance between the two, however, clearly favors portability, with excellent battery life and a slim, semi-rugged design.

Weighing just under 2.5 Kg, the ThinkPad W550s is Lenovo's lightest and thinnest workstation to date, but it's also built to stand up to the rigors of regular business travel. It has a magnesium-alloy roll cage, a carbon-fiber-

reinforced chassis, and stainless steel hinges, and is built to meet MIL-STD-810G standards, meaning it is designed to survive drops, shocks, vibration, and extreme temperatures.

The W550s features an impressive 15-inch display with 3K (2,880-by-1,620 pixels) resolution. With an In-Plane Switching (IPS) panel, the display looks sharp and vibrant from any angle, and it's a 10-digit touch screen, so it supports all of the gesture controls found in Windows 8.

FEATURES

The ThinkPad W550s is well-equipped. On the right of the system are two USB 3.0 ports and a mini DisplayPort. The left side sports a third USB 3.0 port, VGA output, a Gigabit Ethernet port, headset jack, and an SD card slot. There is also an assortment of security and management features, ranging from a case-lock slot for physically securing the system to a built-

in fingerprint reader for simple secure login. A docking connector on the underside of the system lets you dock it to the ThinkPad Ultra Dock, which provides additional connections for power, monitor, and peripherals. The system also boasts both 802.11ac dual-band Wi-Fi and Bluetooth 4.0, and supports Intel WiDi for wirelessly streaming HD video and audio.

The 512GB solid-state drive (SSD) is double the capacity of the drives found in the HP ZBook 14 \$1,599.79 at Amazon or the Lenovo W540. There also aren't too many preinstalled programs taking up that space. The laptop comes with a free 30-day trial of Microsoft Office 365 for new customers, and a 30-day trial of Norton Internet Security, but otherwise you have one or two apps, like Evernote Touch, Skype, and The Weather Channel. But while there aren't a lot of apps preinstalled, the system is made for broad compatibility, with independent software vendor (ISV) certification for a wide assortment of apps, with uses in engineering, design, and finance.

PERFORMANCE

The ThinkPad W550s is equipped with a dual-core Intel Core i7-5600u processor, an ultra-book-class CPU is used here for its blend of processing power and efficiency. Paired with 16GB of RAM, it does help the system offer excellent battery life, but it also means that the raw performance may not be up to your expectations for a workstation. In PCMark 8 Work Conventional, for example, the ThinkPad W550s scored 2,736 points, lagging behind last year's Lenovo W540 (3,105 points) and the HP ZBook 15u G2 (3,124 points), though it does outperform the Dell Precision M3800 (2,664 points). In Cinebench R15, which takes advantage of multiple processing cores, the system lagged behind, scoring 274 points, while quad-core-based systems like the Dell M3800 (599 points) and the Lenovo W540 (637 points) scored drastically higher.

In 3D rendering and other graphics-intensive tests, the laptop was good, but not great. Lenovo outfitted the system with an Nvidia



Quadro K620M graphics card, which offers the ISV certifications and capabilities necessary for a workstation, but with significantly less muscle than the Nvidia Quadro K2100M used in other systems, like the Lenovo W540 and the HP Zbook 15. Instead, it's closer in capability to the HP ZBook 15u G2, which is better suited to tasks like digital content creation and basic financial number crunching than traditional workstation uses like engineering and design.

The ThinkPad W550s has two batteries, an internal battery and a removable secondary one. In our rundown test, the system lasted 6 hours 44 minutes using the internal battery alone, outlasting both the HP ZBook 15u G2 and the Lenovo W540 (both 6:13) by a half-hour. With the addition of the second battery, the time is extended by more than 10 hours (17:21).

CONCLUSION

The Lenovo ThinkPad W550s is a solid laptop for the user who needs basic workstation capabilities, but puts portability first. If you are frequently on the road, the ThinkPad W550s is a great choice, thanks to the thin and light design, combined with the extremely impressive 17+ hours of battery life. In terms of performance, however, the Lenovo ThinkPad W540 is still ahead of the game, offering more processing power and better graphics capability at the expense of portability.



HPC ACADEMY

Developing High-Performance Computing Expertise in Europe



NONPROFIT ORGANIZATION

Our missions:

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

Our first partner members:



Become a member
www.hpc.academy

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

Starting to use the second PFLOPS supercomputer at the Centre for Computational Sciences, University of Tsukuba Center for Computational Sciences, University of Tsukuba (hereinafter, Center for Computational Sciences), promoting “Interdisciplinary Computational Science”, a fusion of science and computer science, introduced “COMA (PACS-IX)” (COMA, PACS-Nine), a supercomputer of peak computing performance of 1.001PFLOPS, and has been in use since April 2014. The CPU comes with two units of Intel Xeon Processor E5 family, the arithmetic accelerator which assists the CPU, combined with 393 units of compute nodes that comes with two units of Intel Xeon Phi coprocessor, form an Ultra parallel PC cluster. It is a matter of pride to achieve the highest

domestic performance (as of June 2014), with peak computing performance of 1.001PFLOPS, with supercomputers having Intel Xeon Phi coprocessors. At the Center for Computational Sciences, in addition to promoting state-of-the-art computational science research, the two supercomputers, ‘HAPACS’, the supercomputer that came into use in 2012 with 1.116 PFLOPS that accelerates the computing device GPU, and the new ‘COMA’ are used towards the professional development of high-performance computing in the education of students and graduate students.

SELF-PROMOTION OF SCIENTIFIC RESEARCH BY USING A UNIQUE COMPUTER AT THE NATIONAL UNIVERSITY INFORMATION CENTRE

With the Centre for Computational Physics that was founded in 1992 as its predecessor, the current Center for Computational Sciences had its restructuring and expansion done



“A LARGE NUMBER OF CPU CORES ARE MOUNTED ON A SINGLE CHIP, AND ARE ATTACHED TO THE CPU VIA THE PCI EXPRESS BUS IN THE INTEL XEON PHI COPROCESSOR TO ACHIEVE HIGH COST PERFORMANCE, POWER PERFORMANCE AND SPACE PERFORMANCE ON THE SUPER COMPUTER”

Assistant Center Head Professor Taisuke Boku

in 2004. In 2010, it was certified by the Joint Usage/Research Center, Ministry of Education as 'AISCI (Advanced Interdisciplinary Computational Science Collaboration Initiative)'.

Currently, under the 'Multidisciplinary Joint-use Program' that strengthens the collaboration of computer science and science, it provides the computer center to research institutes across the country, supports holding of workshops for interdisciplinary computational science promotion, supports invitation to researchers, supports exchange of students and researchers and conducting of research aimed at the development of extra scale supercomputers.

Deputy Director and Chairman of the Computer System operation, Mr. Taisuke Boku, explains thus: “Its role is not just that of the information center where researchers from the field of science can use supercomputers but being a unique National University with a large number of teachers and researchers from various computational science fields, application development and scientific research using this computer, and promoting the development of the computer itself, a feature that is not present in the computer centers of other universities. These initiatives are defined as 'Interdisciplinary Computational Sciences', and as a major feature of the Center for Computational Sciences it is a place of exploration in computational sciences that have a wide scope.” 'Interdisciplinary Computational Science' uses a high-speed network and ultra-highspeed computers as the main research tool, where researchers from three fields, viz. science researchers who use the computer, information science researchers



CHALLENGES

- Expanding needs of the high-speed system to support 'Interdisciplinary Computational Science',

a fusion of science and computer science.

- Accumulation of technology of many-core processors

SOLUTION

- Intel Xeon processor E5 family

- Intel Xeon Phi coprocessor

- Comprehensive tool 'Intel Cluster Studio XE' towards MPI developers and C++/ Fortran programmers

BENEFITS

- Introduction of supercomputer 'COMA (PACS-IX)' of peak computing performance of 1.001PFLOPS

- Research and development of next generation supercomputers that employ many-core processors



THE JOINT CENTER FOR ADVANCED HIGH PERFORMANCE COMPUTING



performance through freedom from vendor dependence, and they are formulated based on the doctrine of “Openness” which represents openness of hardware and openness from user dependence. The establishment of JCAHPC is one step in this direction, and the design and development as well as the operation and management of the high-performance super computers are jointly carried out by the information centers of both universities with the objective of promoting state-of-the-art computational science.

The Center for Computational Sciences, University of Tsukuba and the Information Technology Center, University of Tokyo have set up a super computer system to be designed primarily by the teaching staff of both institutions in the Tokyo University Information Technology Center located in the Kashiwa Campus of the University of Tokyo used since April 2015, and have commenced activities for the launch of a new super computer system. This is the first time such an experiment has been conducted within Japan to set up such a facility and to jointly operate and manage a super computer.

3 CENTERS FOR HPC ACADEMIC RESEARCH

The Center for Computational Sciences, University of Tsukuba and the Information Technology Center, University of Tokyo have implemented projects in the past to set joint specifications and procure super computers for each university through the “T2K Open Super Computer Alliance” consisting of 3 centers of the Academic Information Media Center of the University of Tokyo. The specifications of the T2K Open Super Computer Alliance are intended to reduce procurement cost and improve system

NEW MANYCORE POLICY

The policy framework which has been established for the design and development of super computing systems does not follow existing super computing products, but instead, uses many-core processors to design a state-of-the-art system, and also links the system to the OS, programming language, numerical computing library and other systems which form the core technology for the software during the design and development process.

who research on data and media processing, and computer science researchers who research on hardware, software, algorithms, and programming use these computers for their integrated but innovative applied research. The Computational Science Research Center, along with building a joint research system over the years, has also been promoting basic science,

high-speed simulation, largescale data analysis, and applied research in information technology. Currently, using 'HA-PACS' and 'COMA', that brought into reality the PetaFLOPS, we aim for a breakthrough in the research and development of computational sciences in the fields of particle physics, space, nuclear physics, material, science, life, and global environment.



EVALUATING THE HIGH COST PERFORMANCE AND THE POWER PERFORMANCE AND ADOPTING OF INTEL XEON PHI COPROCESSOR

In the recent years, speed and performance of CPU have helped performance of massively parallel PC clusters to increase significantly. However, the conventional way of just increasing the number of nodes by stacking CPU to increase computing performance has limitations due to limited space and power. At the Computational Science Research Center, a course concerning the study of low power high performance supercomputers accelerated by arithmetic devices that assist CPU focuses on many-core processors (MIC) in which multiple cores are integrated in a single chip. Intel has been participating in the 'Intel MIC Beta Program' initiated for evaluation of the architecture of MIC and continuously researching on tuning and performance evaluation. The research center has decided to introduce COMA instantaneously after the Intel Xeon Phi coprocessor is launched.

The Intel Xeon Phi coprocessor has 61 CPU cores on a single chip. Just attaching this coprocessor to a CPU through PCI Express (generic bus) will achieve a high cost performance as well as bring out the performance of teraflops. This coprocessor gives an excellent space performance and high power performance for every 1 watt. (Mr. Boku) Another benefit of the Intel Xeon Phi coprocessor, which shares same architecture as that of the Intel Xeon processor, is its high level of programmability. . Mr. Boku says, "There comes a sense of security when programs can be developed without wasting the existing programs and without having to remember new things. Just like the conventional coprocessors, this coprocessor can be used to write programs by using FORTRAN or C++ and Open MP can be used, so that developing programs for the future is possible" In 2013, Computational Science Research Center had designed the next supercomputer system in collaboration with the information Technology Center of the University of Tokyo and has crea-

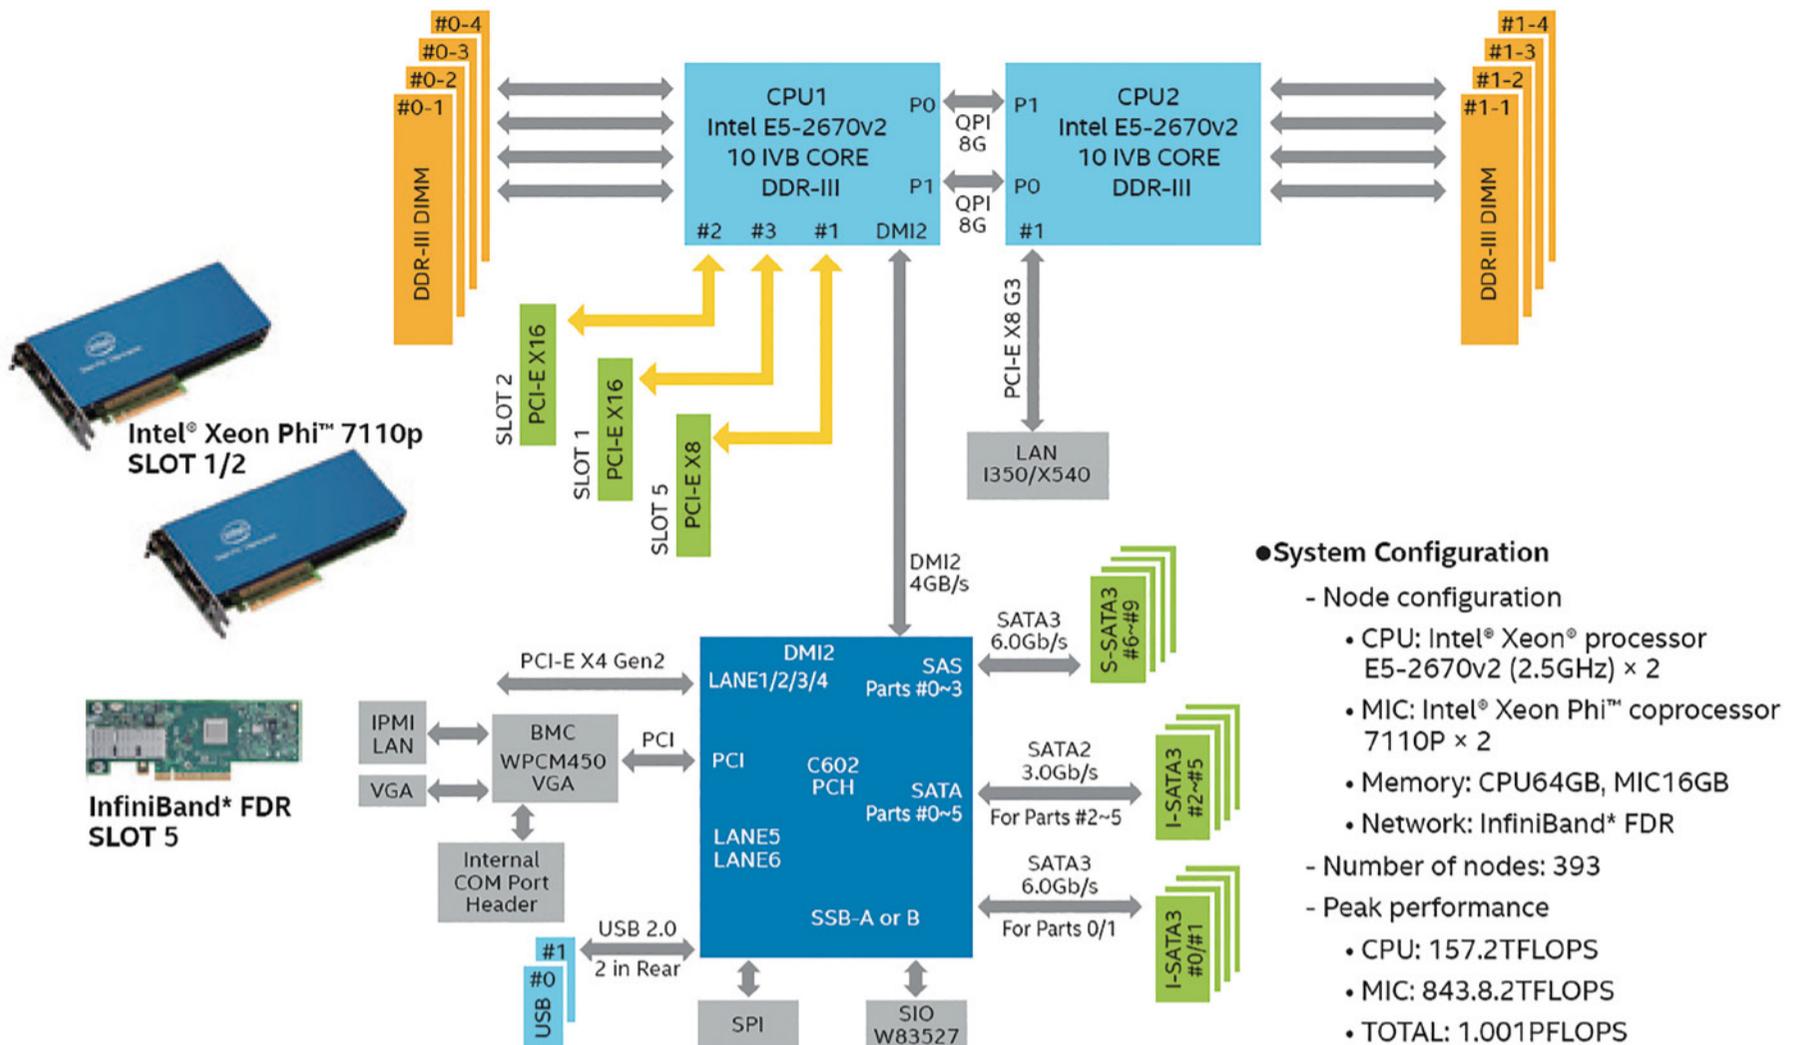
ted 'JCAHPC (Joint Center for Advanced High Performance Computing)' an organization for co-operation and management. As the goal for FY 2015, JCAHPC is planning to introduce very large supercomputers with scores of PFLOPS class, at the Kashiwa campus of the Information Technology Center, University of Tokyo. COMA, which is provided with many-core processors, also plays a role as an experimental system for carrying out the research and development of tuning and coding of the next generation supercomputer.

ENSURING HIGH-SPEED PERFORMANCE WITH BIASED MIC BOARD OF 2 UNITS WITH ONE CPUMIC AND CPU

COMA, introduced by the Center for Computational Sciences is an abbreviation of 'Cluster of Many-core Architecture processor' The term is also derived from the English name of 'Coma Berenices', one of the typical clusters of galaxies. Collection of stars is a Galaxy (Many Core), and a collection of Galaxies (Cluster) has the image of a Galaxy cluster. At the same time, from the fact that it is a machine of the 9th generation Supercomputer PACS Series of the Center of Computational Sciences, the code name of PACS-IX has been used in combination. COMA is a parallel system equipped with 393 compute nodes where one compute node has two Intel Xeon processors E5-2670v2 (2.50GHz) having 10 cores per CPU and two Intel Xeon Phi™ coprocessors 7110P with 61 CPU cores. Peak computing performance of a single node is 0.4TFLOPS for CPU, 2.147TFLOPS for MIC, and 2.547TFLOPS together. Peak computing performance of the entire system is 157.2TFLOPS for CPU, 843.8TFLOPS for MIC, and 1.001PFLOPS together. In addition, all compute nodes are connected by a mutual network through a high-speed network (InfiniBand FDR), with a communication performance ten times that of Ethernet and run a uniform parallel processing between each compute node.

Lustre file server of total capacity 1.5PB by RAID-6 is used for storage and through the InfiniBand* FDR, it is possible to freely access from all compute nodes. "Compute nodes consists of

SYSTEM BLOCK DIAGRAM



two MIC boards equipped with the Intel Xeon Phi coprocessor 7110P and a network board of InfiniBand FDR, next to one Intel Xeon processor E5-2670v2, ensuring high-speed communication capability between the coprocessor and the network board. The merit is that just one Intel Xeon processor E5-2670v2 alone can maintain a high-speed performance.” (Mr. Boku)

USES 3 OPERATING MODES TO SUPPORT PARALLEL PROCESSING OF COMPUTATIONAL NODES

COMA operations can handle three jobs viz. CPU only, MIC only, and a combination of CPU and MIC and are intended to be used in the following way.

• **CPU ONLY** To support the multi-core oriented programs (+MPI) of the previous generation super computer “2K-Tsukuba”, implementation of which has been stopped in February 2014. Uses 16 cores of the total of 20 cores in the Intel Xeon processor E5-2670v2.

• **MIC ONLY** To use four cores of the Intel Xeon processor E5-2670v2 and two cores of the Intel Xeon Phi coprocessor 7110P (61 core × 2). Using the offload functionality provided by the Intel compiler for offload execution of some operations on MIC.

• **CPU+MIC** Monopolization of all computational resources (CPU+MIC) at the node level. To use hybrid programs in the MIC Native mode or Symmetric mode. Intel Linux version of the development tool “Intel Cluster Studio XE” is used in the programming environment, and researchers use the Intel FORTRAN compiler/ Intel C++ compiler, and the Intel MPI library as per their needs to create the programs. Mr. Boku says, “Since free compilers can also be effective for some programs, we have developed an application system in which we do not specify a particular product for the compiler or library, but instead, allow the researchers to freely choose and use a compiler or library of their choice”. Used for development



of computational science applications in fields such as particle theory and life science COMA can be used by researchers throughout the country without any charge under the “Interdisciplinary joint utilization program” to support advanced scientific research in the fields of computational science and computational engineering. And, arrangements have also been made for its free use under the “HPCI program” promoted by the Ministry of Education, Culture, Sports, Science and Technology, and under the “Large scale general use program” which is used by researchers throughout the country free of cost. Research scholars belonging to the Center for Computational Sciences are carrying out research by using “COMA”’s high speed parallel computing in the development of computational science applications in various fields such as the field of Particle Theory, the fields of Astrophysics, Life Sciences and Physical Sciences, and the fields of Material Sciences, Global Environment and Computational Information. Professor Boku has this to say about where to use the large scale parallel GPU cluster “HA-PACS” which first began operations in 2012, and where to use COMA.

“There is no classification based on the subject matter of the research. The researchers can choose as they like, and after careful examination of the subject matter of the research, they can make use of the interdisciplinary joint utilization program. But, HA-PACS require programs to be coded for GPU, as a result of which a large number of programs which run on it have been optimized for GPU. In contrast to this, COMA which uses Intel architecture is highly generic, and can run many more programs.”(Professor Boku)

By delivering lectures on high-performance computing to graduate school students, we develop competence in them to carry out computations according to the techniques of computational science COMA is presently also active in the education of students of the Center for Computational Sciences which promotes interdisciplinary computational science. In the words of Professor Boku, “Today, the reality is that

Japan is running overwhelmingly short of researchers in the scientific field who can write their own programs, when compared to the situation abroad. Therefore, in order to develop competence in many of our researchers so that they can write their own computational programs in various fields of research to perform computations according to techniques of computational science, we are conducting classes in high-performance computation literacy for graduate school students. Besides this, we are also running intensive courses on introduction to high performance computing for 40 to 50 graduate school students every year, and continuing with our initiatives to send out qualified persons with knowledge of high-performance computing”

And even in our “Global 30 (G30)” study program for overseas students who can only attend classes and obtain credits in English, we are planning to conduct classes on super computers such as COMA and HA-PACS directed at these overseas students in order to offer advanced educational services. In the Center for Computational Sciences, we plan to continue to use COMA in various fields of research and study the performance of many-core processors in order to gain knowhow and knowledge to help in the achievement of high-performance. In addition, we are planning to continue with research and performance evaluation in order to achieve the planned launch in 2105 of the super computer using a many-core processor to be jointly operated by the Center for Computational Sciences and the Information Technology Center, University of Tokyo.

Finally, Professor Boku says about Intel that “To implement EFLOPS, they are waiting for feedback based on verification and evaluation of super computing technology in various fields”. Intel will continue to support the development of inter-disciplinary computational science which is being promoted by the Center for Computational Sciences by improving the performance of the Intel Xeon processor and Intel Xeon Phi coprocessor, and making advancements in MIC architecture.



HPC ACADEMY

Developing High-Performance Computing Expertise in Europe

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

NEXT HPC TRAINING

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

About the training:

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

Target audience:

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

Language: French

Duration: 5 days

Location: Laval, France

Price: 4 000 euros

For more informations



Next session:

November 2nd to November 6th 2015

Our first partner members:



Become a member
www.hpc.academy



HOW LUSTRE ADDS VALUE FOR COMMERCIAL HPC

A

growing number of industries require high-performance computing (HPC) clusters to enable discoveries hidden in the massive amounts of data they collect. In the field of genomics, for example, researchers at Iowa State University work with hundreds of gigabytes for studies of a single genome.

These scientists, working on assembling sequences of *Zea mays*, are dealing with multi-gigabyte genomes, but they must sequence a single genome 150X to get reliable data from it for their research. In another example, for climate models used in weather prediction, the high resolutions of today's simulations can generate as much as 100 terabytes of data, according to a "Journal of Advances in Modeling Earth Systems" report by Michael F. Wehner and others. Analyses from high resolution seismic data similarly can rely on incredibly large data sets of over 100 terabytes in a single file.

To reliably process this magnitude of data in a reasonable amount of time using HPC, very few storage architectures can keep up with the computing systems enterprises and academia use. The solution is usually to deploy a parallel

file system (PFS) with I/O capable of hundreds of gigabytes to several terabytes/second. The majority of parallel file systems being deployed today are IBM's Spectrum Scale* (formerly called General Parallel File System*, or GPFS*) and open source Lustre*. These two solutions present completely different purchasing models for organizations looking to deploy a high-performance data platform for massively large data sets. Lustre, being a community-developed and open source licensed model, can create significant value for high-performance data solutions.

A LITTLE PFS HISTORY

Spectrum Scale (or GPFS) was developed at IBM in the early 1990s and finally commercialized at the end of that decade. It can be deployed in different distributed parallel modes and is used in large commercial and supercomputing applications. It began its life on IBM's AIX systems, but has since appeared on Linux* and Windows Server* systems.

Lustre's architecture began as an academic project at Carnegie Mellon University in the late 1990s. The file system was developed under the Accelerated Strategic Computing Initiative (ASCI) as part of a project funded by the U.S. Department of Energy that included Hewlett-Packard and Intel. (ASCI drove the creation of the first teraflop supercomputer, the Intel-based ASCI Red installed at Sandia

National Labs in 1996.) Today, Lustre development continues under an open source model, with releases managed by OpenSFS*. Intel Corporation contributes the majority of code enhancements to Lustre, while also adding its own set of unique features in several Intel-branded versions that make it more desirable for enterprise and cloud applications and deployments.

PARTNERSHIPS FOR SUCCESS

High-performance storage solutions capable of serving today's petascale and tomorrow's exascale computing are complex and large systems with tens of thousands of spinning or solid-state drives. These are not typically something an in-house IT department will tackle alone. Organizations usually seek out a partner with expertise in the software, hardware, and overall requirements for scalability, optimization, configuration, networking, backup, and disaster recovery. And therein lies one of the critical decisions companies face in finding their high-performance data solution—to go with a closed-source, fee-based licensing solution with potential vendor-locked-in or expand their choices and flexibility for the system makeup that incorporates both hardware neutrality, so purchasers can choose from a wide variety of manufacturers and technologies, and an open source licensing model with Lustre.

Many customers choose a proprietary-based solution from enterprises that have strong reputations of expertise, reliability, and trust. But, in today's marketplace, dollars are short and competition is fierce. And choice is a key ingredient that makes the economics of technology efficient, successful, and cost-effective. New company entrants in high-performance data solutions designed around Lustre have been making serious inroads to the enterprise storage market for several years.

There is a wide selection of smaller yet equally competent technology integrators around the world who have developed storage solution expertise with Lustre, while deliv-

ring a level of personal attention and instant response to customers that rivals the big and expensive enterprises. Most of these integrators are part of Intel Corporation's Lustre reseller program, which numbers more than 20 companies. They can offer highly competitive solutions around Lustre.

ADDING VALUE THROUGH OPEN SOURCE

Open source has been the mode of enterprises deploying cost-effective solutions built on reliable, performant software for decades. Linux is the open source poster child, and Red Hat is the standard bearer for building a profitable enterprise business around Linux by providing enterprise-grade services along with its own productized enhanced offerings around the Linux kernel, while continuing to contribute to the community's development tree. Other companies, like Novell with their SUSE distribution, have followed the same model.

When Lustre entered the open source community, several organizations began to form that wanted to maintain Lustre's traction and momentum in HPC. They also grasped the incredible potential of Lustre for both academia and enterprise. So, they began building on the same open source model: contributing to the code while offering enhanced product and support offerings around the software. This was good for enterprise, because, similar to the expansion of Linux into the enterprise, company IT managers wanted the assurance of support and longevity behind Lustre before committing critical applications to the file system.

"Intel is the clear leader in the Lustre Open Source project," according to Brent Gorda, Intel's General Manager of their High Performance Data Division (HPDD), the group responsible for Lustre at the company. "We are the chief contributors to the Lustre code and have the largest pool of Lustre experts in the community. Building on the open source version, we offer enterprise-class and cloud versions of Lustre with Intel software enhancements and enterprise-grade services to cus-

tomers. In this way we are adding value to open source Lustre and helping commercial HPC easily take advantage of the fastest scalable parallel file system on the planet.”

According to Gorda, Lustre eliminates vendor lock-in and enables more choices to customers. If they choose to use a technology partner for their high-performance data solution, they can RFP across many vendors to obtain the optimum hardware technology at the best value, and use Lustre as the software. This way, they avoid the often steep and changing license fees that go along with proprietary solutions, while still getting the support expertise they need.

LUSTRE IN LEADING INSTALLATIONS

Lustre is the high-performance file system of choice for many large installations, including Lawrence Livermore National Laboratory’s Sequoia supercomputer. Core to the HPC clusters at San Diego Computing Center (SDSC) at the University of California, San Diego is Data Oasis, a scalable, Lustre parallel file system with up to 12 PB (petabytes) of capacity and exceeding 200 gigabytes/second while supporting 10,000 simultaneous users. Designed and deployed by Aeon Computing, also of San Diego, the 72-node system is linked to SDSC’s Trestles, Gordon, TSCC, and Comet clusters. Data Oasis’ sustained speeds mean researchers can retrieve or store 240 TB of data in about 20 minutes. Early this year, Data Oasis began undergoing significant upgrades, including ZFS, a combined file system originally designed by Sun Microsystems and mated in a new hardware server configuration under a partnership between SDSC, Aeon Computing, and Intel.

GAME CHANGING OFFERINGS

As the core to today’s Open Source, high-performance, scalable data storage solutions, Lustre offerings, with Intel enhancements and support, present a valuable alternative to closed-source, proprietary solutions. Thus, Lustre has been a game changer for high-per-

formance data storage systems in academia and enterprise HPC. And Lustre continues to undergo critical enhancements that enterprise wants, with releases managed through OpenSFS and the Lustre community.

Intel’s Lustre offerings include Intel® Enterprise Edition for Lustre software, Intel® Cloud Edition for Lustre software, and Intel® Foundation Edition for Lustre software. Each is designed to meet specific needs in the marketplace. The Enterprise Edition brings together powerful tools and features in the latest Lustre release that enterprise IT has been requesting to make it easy to deploy, to significantly expand the number of metadata servers, increase reliability, include hierarchical storage management, and allow it to interface with Hadoop workloads. Support is available directly and through Intel resellers, giving companies greater choice without sacrificing the ease-of-manageability and support they want in a storage solution. The Cloud Edition, available through the Amazon Web Services Marketplace, allows customers to stand up a scalable parallel file system in minutes for any number of applications a customer wants to run on Amazon’s Elastic Compute Cloud (EC2). This open source approach from the Lustre community and from companies that are building operations around the software, add considerable value to customers needing high-performance storage solutions today.

ABOUT THE AUTHOR

Ken Strandberg is a technical story teller. He writes articles, white papers, seminars, web-based training, video and animation scripts, and technical marketing and interactive collateral for emerging technology companies, Fortune 100 enterprises, and multi-national corporations. Mr. Strandberg’s technology areas include Software, HPC, Industrial Technologies, DesignAutomation, Networking, Medical Technologies, Semiconductor, and Telecom. Mr. Strandberg can be reached at ken@catlowcommunications.com



FAST
SSDS FOR
BIG DATA
PROCESSING
REQUIRES
HIGH-
PERFORMANCE
STORAGE

IT IS IMPORTANT TO UNDERSTAND THE WORKLOAD PERFORMANCE AND ENDURANCE REQUIREMENTS BEFORE MAKING A DECISION.

Big data applications handle extremely large datasets that present challenges of scale. High-performance IT infrastructure is necessary to achieve very fast processing throughput for big data. Solid state drives (SSDs) based on NAND Flash memory are well-suited for big data applications because they provide ultra-fast storage performance, quickly delivering an impressive return on investment. SSDs can be deployed as host cache, network cache, all-SSD storage arrays, or hybrid storage arrays with an SSD tier.

Depending on the big data application, either enterprise-class or personal storage SSDs may be used. Enterprise SSDs are robust and durable, offering superior performance for mixed read/write workloads, while personal storage SSDs typically cost less and are suitable for read-centric workloads. It is important to understand the workload performance and endurance requirements before making a decision.

THE BIG PICTURE OF BIG DATA

In a scenario where data grows 40% year-over-year, where 90% of the world's data was created in the last two years, and where terabytes (TBs) and petabytes (PBs) are talked about as glibly as megabytes and gigabytes, it is easy to mistakenly think that all data is big data. In fact, big data refers to datasets so large they are beyond the ability of traditional database management systems and data processing applications to capture and process. While the exact amount of data that qualifies as big is debatable, it generally ranges from tens of TBs to multiple PBs. The Gartner Group further characterizes it in terms of volume of data, velocity into and out of a system (e.g., real-time processing of a data stream), and variety of data types and sources. Big data is typically unstructured, or free-floating, and not part of a relational database scheme.

Examples of big data applications include:

- **BUSINESS ANALYTICS** to drive insight, innovation, and predictions
- **SCIENTIFIC COMPUTING**, such as seismic processing, genomics, and meteorology
- **REAL-TIME PROCESSING** of data streams, such as sensor data or financial transactions

SSDS ARE AN ORDER OF MAGNITUDE DENSER AND LESS EXPENSIVE THAN DRAM, BUT DRAM HAS HIGHER BANDWIDTH AND SIGNIFICANTLY FASTER ACCESS TIMES. COMPARED TO HDDS, SSDS OFFER ORDERS OF MAGNITUDE FASTER RANDOM I/O PERFORMANCE AND LOWER COST PER IOPS, BUT HDDS STILL OFFER THE BEST PRICE PER GIGABYTE.

- **WEB 2.0 PUBLIC CLOUD SERVICES**, such as social networking sites, search engines, video sharing, and hosted services

The primary reason for implementing big data solutions is productivity and competitive advantage. If analyzing customer data opens up new, high-growth market segments; or if analyzing product data leads to valuable new features and innovations; or if analyzing seismic images pinpoints the most productive places to drill for oil and gas—then big data is ultimately about big success. Big data presents challenges of extreme scale. It pushes the limits of IT applications and infrastructure for processing large datasets quickly and cost-effectively. Many technologies and techniques have been developed to meet these challenges, such as distributed computing, massive parallel processing (e.g., Apache Hadoop), and data structures that limit the data required for queries (e.g., bitmaps and column-oriented databases). Underlying all of this is the constant need for faster hardware with greater capacity because big data requires fast processing throughput, which means faster, multicore CPUs, greater memory performance and capacity, improved network bandwidth, and higher storage capacity and throughput.

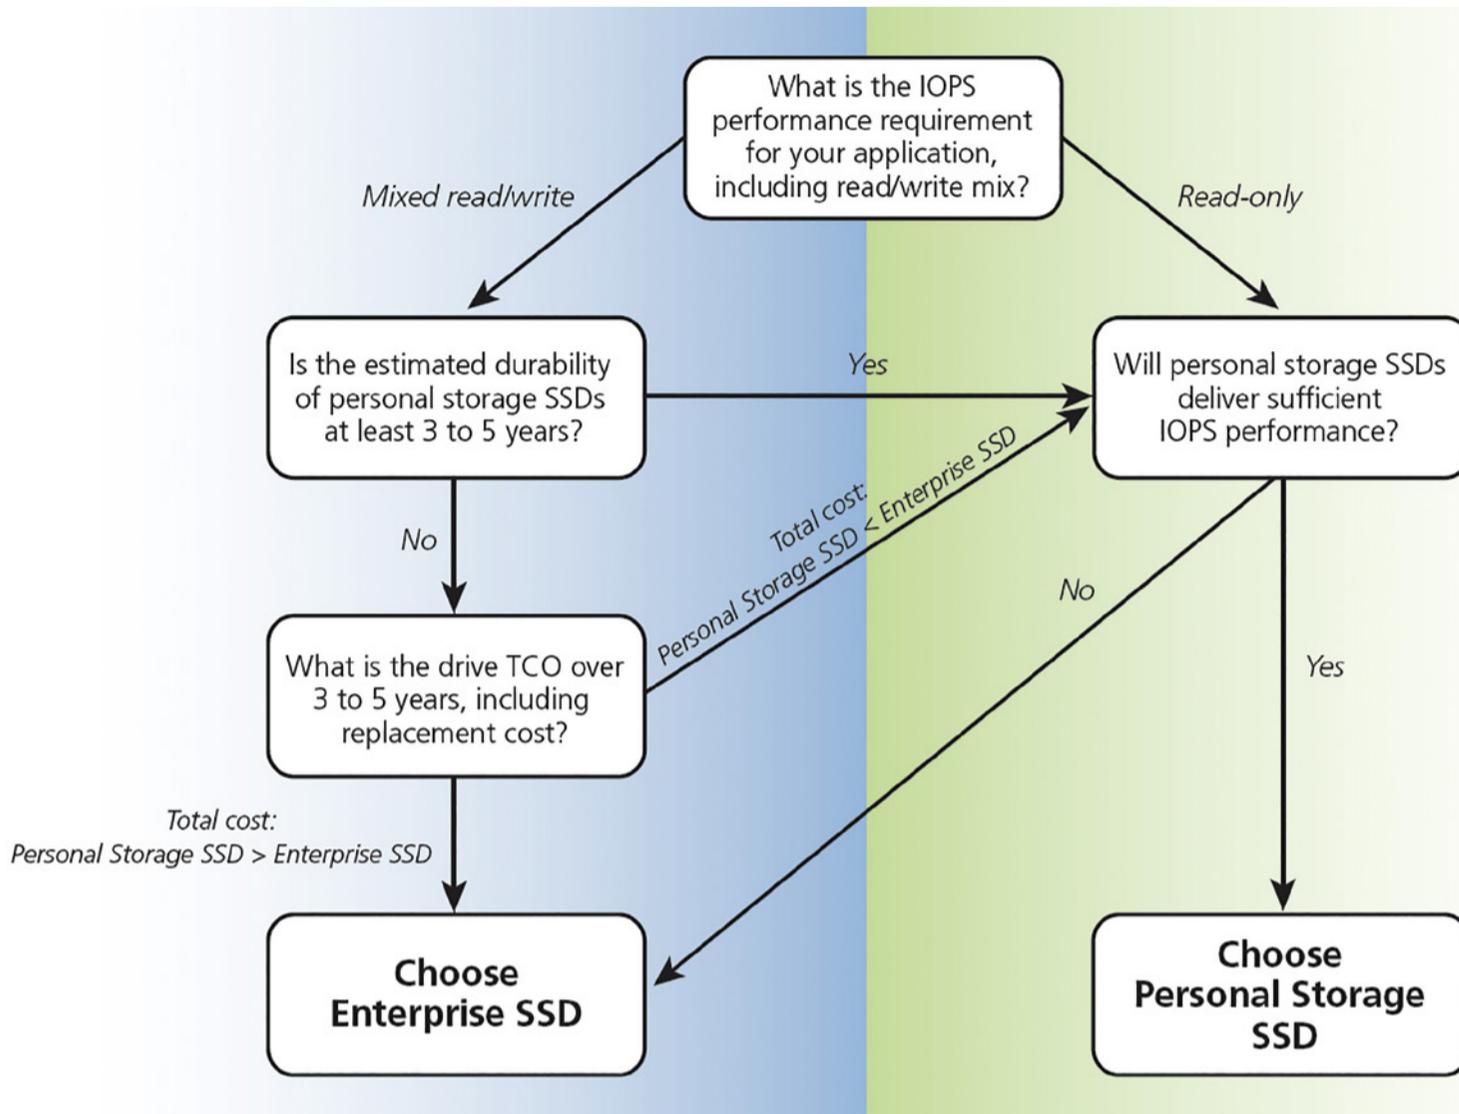
SSDS FOR ULTRA-FAST STORAGE

SSDs have emerged as a popular choice for ultra-fast storage in enterprise environments, including big data applications. SSDs offer a level of price-to-performance somewhere between DRAM and hard disk drives (HDDs). SSDs are an order of magnitude denser and less expensive than DRAM, but DRAM has higher bandwidth and significantly faster access times. Compared to HDDs, SSDs offer orders of magnitude faster random I/O performance and lower cost per IOPS, but HDDs still offer the best price per gigabyte. With capacity pricing for Flash memory projected to fall faster than other media, the SSD value proposition will continue to strengthen in the future.

SSD Benefits

- **EXCEPTIONAL STORAGE PERFORMANCE** – Deliver good sequential I/O and outstanding random I/O performance. For many systems, storage I/O acts as a bottleneck, while powerful, multicore CPUs sit idle waiting for data to process. SSDs remove the bottleneck and unleash application performance, enabling true processing throughput and user productivity.
- **NONVOLATILE** – Retain data when power is removed; no destaging required, like DRAM.

FIGURE 1. DECISION TREE FOR ENTERPRISE-CLASS VS. PERSONAL STORAGE SSD



- **LOW POWER** – Consume less power per system than equivalent spinning disks, reducing data center power and cooling expenses.
- **FLEXIBLE DEPLOYMENT** – Available in a unique variety of form factors and interfaces compared to other storage solutions:
 - **FORM FACTORS** – Half-height, half-length (HHHL), 2.5-inch, 1.8-inch, mSATA, m.2, etc.
 - **INTERFACES** – PCIe, SAS, and SATA

SSD Deployment Options

- **HOST CACHE** – SSDs reside in the host server and act as a level-2 cache for data moved out of memory. Intelligent caching software determines which blocks of data to hold in cache. Typically, PCIe SSDs are used because they offer the lowest latency because no host controllers or adapters are involved. Best results are achieved for heavy read workloads. Cache may be read-only or write-back. Redun-

Drive TCO = Cost of drives + Cost of downtime + Cost of slowdown + Cost of IT labor + Risk of data loss

dant SSDs are recommended for write-back to ensure data is protected.

- **NETWORK CACHE** – Similar to host cache, except SSDs reside in a shared network appliance that accelerates all storage systems behind it. Out-of-band cache is read-only, while in-band is write-back. Network cache offers a better economic benefit because it is shared, but it can be slower than direct host cache.
- **ALL-SSD STORAGE ARRAY** – An enterprise storage array that uses Flash for storage and DRAM for ultra-high throughput and low latency. All-SSD arrays offer features like built-in RAID, snapshots, and replication traditionally found in enterprise storage. They may include technologies like inline compression

PERSONAL STORAGE SSDS ARE DESIGNED FOR GOOD READ PERFORMANCE AND TAILORED RELIABILITY AND DURABILITY. THEY ARE OPTIMIZED FOR WORKLOADS WHERE READS ARE MORE FREQUENT THAN WRITES.

and deduplication to shrink the data footprint and maximize SSD efficiency. This option provides additional management of SSDs as it relates to wearout over the entire array.

- **SSD TIER IN A HYBRID STORAGE ARRAY**

– A traditional enterprise storage array that includes SSDs as an ultra-fast tier in a hybrid storage environment. Automated storage management monitors data usage and places hot data in the SSD tier and cold, or less-frequently accessed, data in high-capacity, slower HDD tiers to optimize storage performance and cost. This option works well for mixed data, some of which requires very high performance. A variation on hybrid storage is when an SSD is incorporated as secondary cache in the storage controller's read/write cache.

CHOOSING THE RIGHT SSD IN BIG DATA DEPLOYMENTS

SSDs in general are rated for 1 or 2 million device hours (MTTF), which translates to at least a century or two of operation. NAND Flash cells only wear out if they are being written to. Enterprise-class SSDs are designed for high reliability, maximum durability, and fast, consistent performance. Enterprise-class SSDs last 10 to 1000 times longer than personal storage SSDs under write workloads. While Flash memory performance tends to degrade with use, enterprise SSDs maintain performance over time. Write performance is 2 to 12 times better, and read performance is comparable to or better than personal storage SSDs.

The price per gigabyte is 2 to 30 times more for enterprise-class SSDs. Big data applications in large corporate data centers, like scientific computing and business analytics applications, are often characterized by mixed read/write workloads that require very low latency and massive IOPS—a good match for durable, robust enterprise-class SSDs.

Personal storage SSDs are designed for good read performance and tailored reliability and durability. They are optimized for workloads where reads are more frequent than writes. Personal storage SSDs offer high capacity and lower price per gigabyte than enterprise-class SSDs. Web 2.0 public cloud applications like social networking sites are characterized by users uploading images, video, and audio files, which are subsequently downloaded or streamed by other users. This type of write-once, read-many-times workload is a good candidate for personal storage SSDs.

APPLICATION CONSIDERATIONS FOR ENTERPRISE VS. PERSONAL STORAGE SSDS

Not all big data deployments are the same, and not all SSDs are the same. The question is how to match the right SSD to the right big data deployment. Choosing an SSD solution is based primarily on the performance and availability requirements of the application. The decision tree in Figure 1 and the following Q&A will help you choose the optimal SSD solution for your application.

THE FIRST STEP IS TO QUANTIFY THE WORKLOAD THAT SSDS WILL SUPPORT.

Q. #1: What is the IOPS performance requirement for your application, including the read/write mix?

The first step is to quantify the workload that SSDs will support. An application workload can be measured using a variety of performance monitoring tools. Beyond workload, also consider the configuration of the system and the impact on the overall platform.

Q. #2: What are the endurance requirements for your application?

For mixed read/write workloads, it is important to look closely at SSD durability ratings. This is usually expressed in terms of total bytes written (TBW) or full drive writes per day over a 5-year period. By comparing an application's daily write total with the durability rating of an SSD, it is possible to estimate the drive's lifetime in your environment (assuming a constant workload; it might be wise to also estimate future growth). If the write workload is small enough that the estimated lifetime of a personal storage SSD will at least equal the 3 to 5 years typically expected of an IT system, and performance is sufficient, then personal storage SSDs can be a good choice. However, if personal storage SSDs will likely wear out and need to be replaced during the IT system's lifetime, then replacement costs should be considered.

Q. #3: What is the SSD total cost of ownership over 3 to 5 years?

The SSD TCO over the system lifetime includes:

- **COST OF DRIVES** – Determine how many drives will need to be purchased during a 3- to 5-year period, including replacements due

to wearout, and multiply this figure by the acquisition cost.

- **COST OF APPLICATION DOWNTIME** – If the application needs to be taken offline to replace an SSD, what is the cost for that lost productivity? Multiply this figure by the number of replacements.

- **COST OF SLOWER APPLICATION PERFORMANCE** – If the application does not have to go offline for drive replacements, but system performance will slow during the replacement and subsequent data replication or RAID rebuild, how will this affect user productivity? Multiply this cost by the number of replacements.

- **COST OF LABOR FOR DRIVE REPLACEMENT** – Drive monitoring and replacement will be an additional management task for the IT staff, so the cost of labor should be included.

- **RISK OF DATA LOSS** – For unprotected drives, there is a significant risk of data loss, and even for RAID-protected drives, there is a small risk during the RAID rebuild window. Though difficult to quantify, these risks should be factored into the cost.

CONCLUSION

SSDs are a popular solution for big data applications. Deciding between personal storage and enterprise-class SSDs will depend on performance and endurance requirements and TCO.

MICRON TECHNOLOGY, INC is a global leader in advanced semiconductor systems. Micron's broad portfolio of high-performance memory technologies—including DRAM, NAND and NOR Flash—is the basis for solid state drives, modules, multichip packages and other system solutions.



HPC ACADEMY

Developing High-Performance Computing Expertise in Europe

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

NEXT HPC TRAINING

HPC02: Developing for performance

About the training:

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

Target audience:

Scientists, technicians and engineers in software development.

Language: French

Duration: 10 days

Location: Laval, France

Price: 8 000 euros

For more informations



Next session:

October 19th to October 30th 2015

Our first partner members:



Become a member
www.hpc.academy

NEXT ISSUE



#8
OCTOBER
31th, 2015

TECH ZONE

**WHY FULL
FLASH STORAGE
IS GAINING
MOMENTUM**

COVER STORY

**BIG DATA
DRIVING THE
WORLD**

LAB REVIEW

**MICROSOFT
ACTIVE
DIRECTORY**



INTERVIEW

SANDEEP GUPTA, SENIOR
DIRECTOR, NVIDIA PROFESSIONAL
SOLUTIONS GROUP

 **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
Renuke Mendis
Marc Paris

ART DIRECTOR
Bertrand Grousset

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