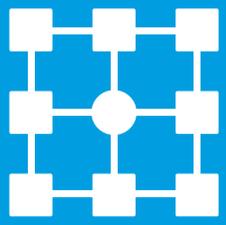


POWER AND GLORY
WORLD'S TOP 10
SUPERCOMPUTERS

7th GENERATION LTO
EXCLUSIVE INTERVIEW
AND 5-YEAR ROADMAP

TECH ZONE
2016-2020 : THE RACE
FOR EXASCALE HAS BEGUN



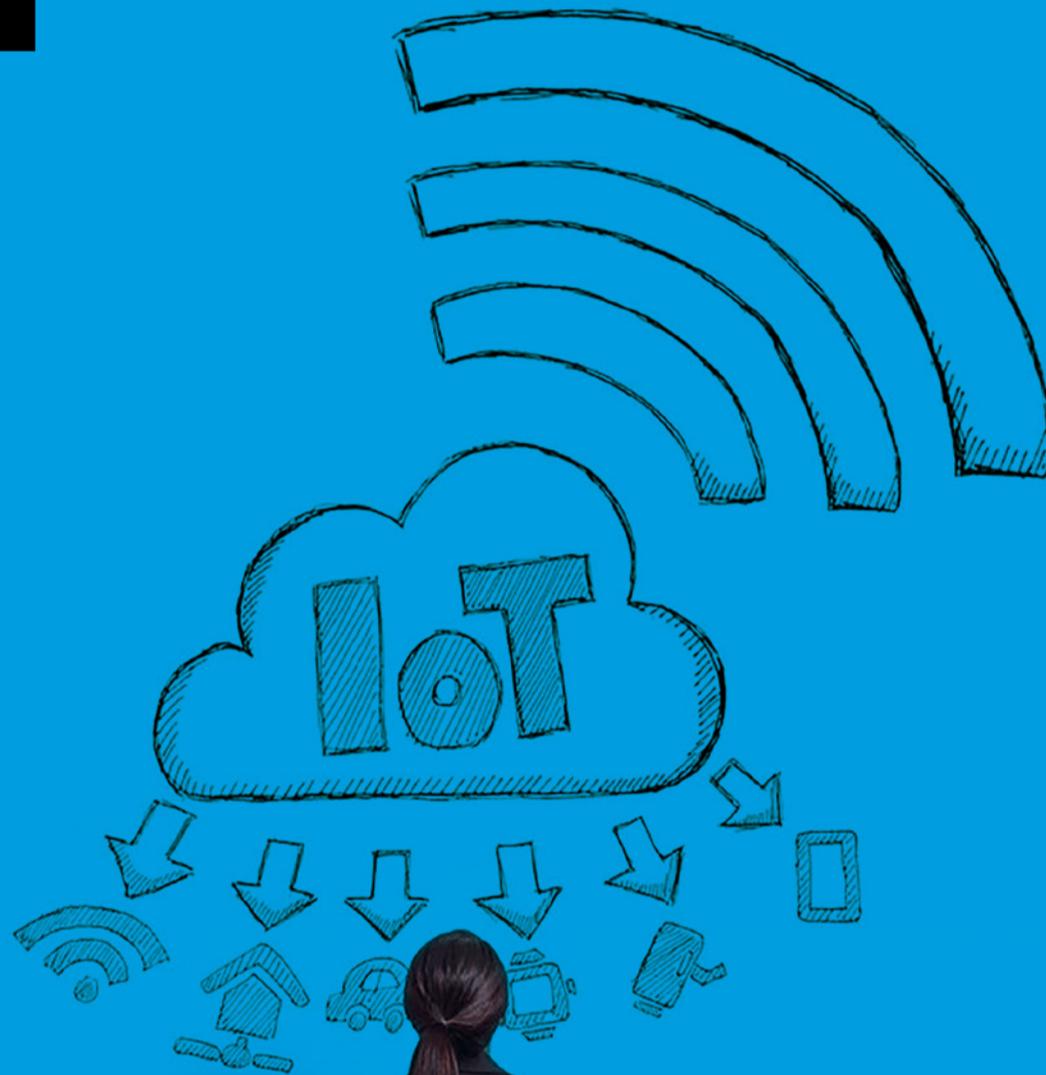
HPC REVIEW

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

www.hpcreview.com

IOT

**THE POWER
TO CHANGE
THE WORLD**



BIRTH OF AN
INDUSTRIAL
ECOSYSTEM

#10 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.

IOT : THE POWER TO TRANSFORM THE WORLD

LET'S FACE IT, VERY FEW ARE THE TECHNOLOGIES THAT HAVE THE POWER TO ACTUALLY CHANGE THE WORLD. Despite the revolutions that the high-tech industry is adamant of trumpeting at every turn, only a handful had this impact over the past century. Among those we can remember the transistor William Shockley invented at Texas Instruments in 1947, a real prelude to integrated circuits that have paved the way for modern computers.

ANOTHER MAJOR ADVANCE WAS MADE UP TO THIS DAY WITH WHAT WAS NOT YET CALLED CLOUD ONLY TEN YEARS AGO. Departmental computers and mainframes began their evolution to the countless datacenters consisting of increasingly dense and powerful servers. A new boundary was crossed.

IN RECENT YEARS WE HAVE SEEN THESE MACHINES GAIN MORE COMPUTING POWER TO ENABLE REAL TIME BIG DATA AND MACHINE LEARNING. Again, this is a prelude, an indispensable foundation for what will follow. Since the era of the Internet of Things, announced for several years now, is beginning to take shape and the industry has begun organizing, structuring, industrializing. It is this decisive step that we present in our cover story.

HAPPY READING!



IOT

**THE POWER
TO CHANGE
THE WORLD**



NEWSFEED

WORLD'S TOP 10 SUPERCOMPUTERS

LTO: 7TH GENERATION IS HERE !

**MIRA: THE WORLD'S FIRST LHC
SIMULATOR SUPERCOMPUTER**

**NVIDIA OPENS THE WAY OF
HYPERSCALE MACHINE LEARNING**

**ASROCK RACK STRATEGY: CUSTOMIZED
AND GREEN HPC SERVERS**

EDITOR'S NOTE
OPINION

SOURCES
THE HPC OBSERVATORY

HOW TO

**IBM WATSON:
THE FUTURE
PROBLEM SOLVING
SUPERCOMPUTER**

VIEWPOINT

**INSIDE INTEL KNIGHTS
LANDING ARCHITECTURE**

TECHZONE

**CERN SEASON 2: NEW
FRONTIERS IN PHYSICS
AND EXASCALE IN 2020**

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.

ENTERPRISE USAGE BOOSTS HPC



EARL C. JOSEPH

IDC RESEARCH VICE PRESIDENT
 AND EXECUTIVE DIRECTOR

Supercomputing 2015 was the opportunity to hear from the server vendors, that the business sector has the highest market growth. This trend has contributed to the growth of the total HPC server market by nearly 12% to \$ 11.4 billion at the end of 2015, according to IDC. Which is an upward revision of the forecasts IDC had at the beginning of the year in part because of the rapid adoption of server computing in the financial services sector. “The combination of HPC and Big Data creates new solutions,” said Earl Joseph, Vice President of HPC at IDC, “with particularly strong growth in the simulation and analysis of data-intensive, and adding many new users and purchasers in the HPC industry. “Among the figures, those who most caught our attention concerned the ROI, return on investment. According to the latest data from IDC, every dollar invested in the HPC sector generated an average 514.7 dollars in revenue. IDC estimates that this growth will be sustained next year if not amplified by the problems related to storage, strong growth also. The overall turnover of the HPC global level should be 22.1 billion. The combination of Big Data and HPC continues to induce changes in the way IDC defines and monitors this emerging market. To qualify, IDC coined the term High Performance

Data Analysis (HPDA). Joseph Earl notes that the HPDA brings out new solutions and added many new users and buyers in the HPC market. The financial sector, for example, “has grown faster than what we measured in the last two years (about 50% higher). “

Currently, IDC has identified four sub-segments for HPDA:

FRAUD DETECTION - This segment focuses on the identification of suspicious or potentially suspicious behavior and their causes using the graphical analysis, semantic analysis, or other high-performance analysis techniques.

MARKETING - This segment covers the use of HPDA to promote products or services, generally using complex algorithms to discern the demographic characteristics of potential customers, their preferences and consumption habits.

BUSINESS INTELLIGENCE - This segment uses HPDA business to identify opportunities to advance their businesses and business competitiveness through a better understanding of their business, their competitors and the changing dynamics of the markets in which they operate.

COMMERCIAL HPDA - This segment includes all commercial HPDA workloads other than the three described above.

THE COMBINATION OF BIG DATA AND HPC CONTINUES TO INDUCE CHANGES IN THE WAY IDC DEFINES AND MONITORS THIS EMERGING MARKET.

Overtime, IDC expects some of these workloads can become important enough to become full-fledged sub-segments. For example, the use of HPDA to manage large IT infrastructures, and the Internet of Things. A future segment HPDA said Joseph, will be precision medicine. A sample application will focus on the analysis of medical outcomes for diagnostic purposes and treatment planning. In this paradigm, history and the patient's symptoms are recorded in a database. While the patient is still in the office of his doctor, the latter will be able to correlate the millions of archived records of patients with similar symptoms to get relevant results. The doctor will be able to take into account the rate of effectiveness of different treatments for patients with similar disease profile, without being bound by the conclusions. In this, the HPDA applied to medicine can be a powerful decision-making tool. Public and private payers have long encouraged the development of similar approaches to evidence-based medicine in which entire populations of patients records could be examined and evaluated to determine which drugs and therapies are useful and should be approved. In theory, the result in both cases would be a better success rate and a reduction in aberrant and costly practices. The cloud-based HPC computing is on the rise in both volume (we went from 13.8% of sites in 2011 to 23.5% in 2013 and 34.1% in 2015), as is the number of running workloads.

Other findings of the IDC report are:

- The growing recognition of the strategic value of HPC helps boost premium sales. The entry-level buyers are also back in a growth mode.

- Marketshare positions of HPC providers will change significantly in 2015.
- Recognition of the strategic /economic value of HPC stimulate the race to exascale, with systems capable of 100 PetaFLOPS power becoming commonplace in 2016. exascale infrastructure should arrive between 2022 and 2024.
- Non-x86 processors could change the landscape – Power processors, but also ARM coprocessors, GPUs and FPGAs.
- China will occupy a more important place. Lenovo, benefiting from the growth of the domestic market, will confirm its export intentions. Other Chinese suppliers are planning to expand to Europe.
- The growing influence of the data center in the IT channel will impact technology options open to the HPC ecosystem, perhaps by providing new approaches.
- The HPC in the Cloud is gaining ground and the big questions are how far and how fast.

Of course there is still room for improvement. The software remains a major obstacle. Improved management software and HPC resource administration is essential and there is a persistent lack of parallelized software for the majority of users. In addition, many applications would need an overhaul or a major rewrite in order to function in HPC environments. Clusters are still notoriously difficult to deploy, use and manage. The issues of power, cooling and floor space remain major challenges. Moreover, IDC indicates that HPC still suffers from a lack of support in heterogeneous environments combining processors and accelerators. Finally, storage and data management are growing bottlenecks given the foreseeable significant increase in data volume.

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)

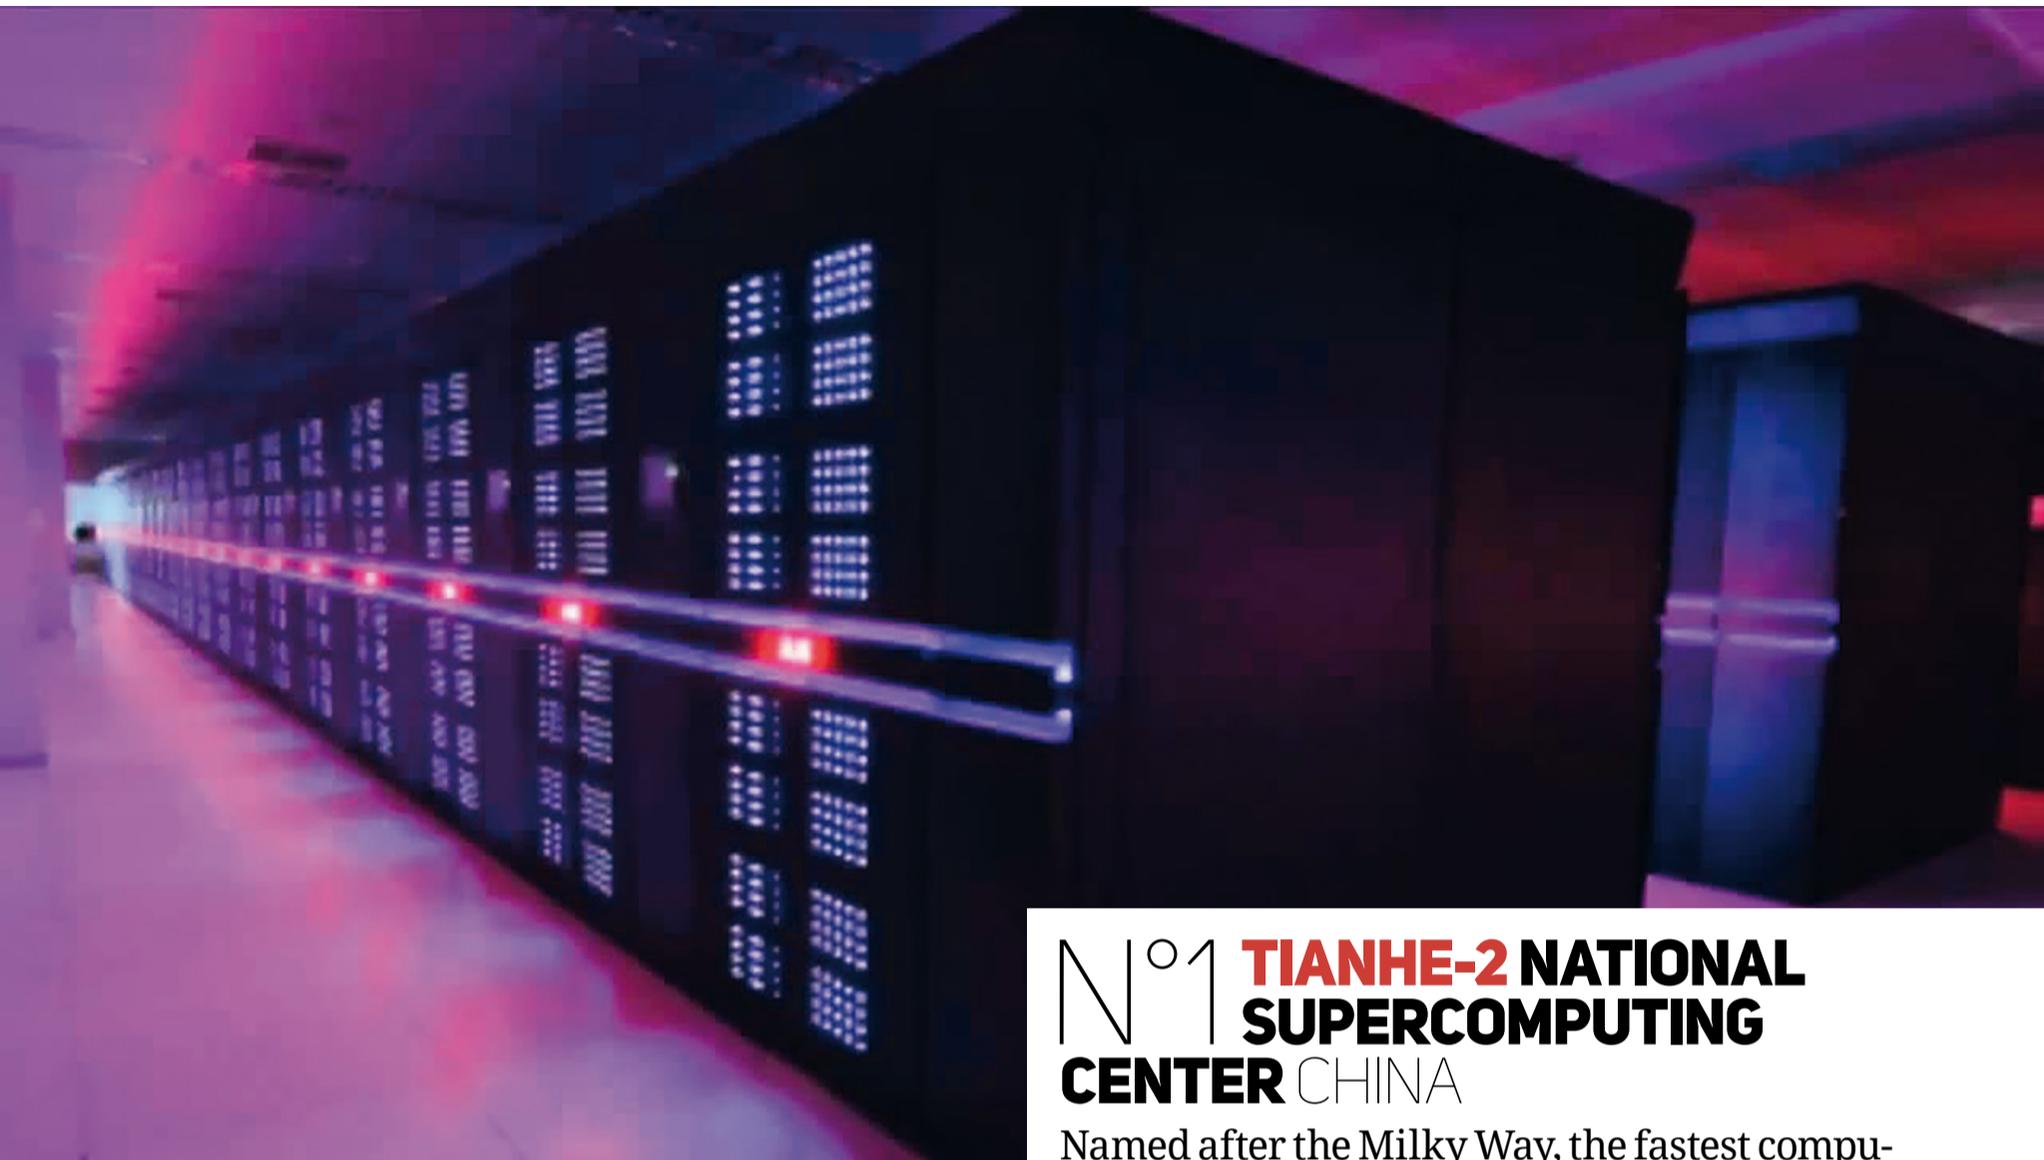
● **NEWSFEED**



POWER AND GLORY
WORLD'S TOP



SUPERCOMPUTERS



N^o1 **TIANHE-2 NATIONAL SUPERCOMPUTING CENTER** CHINA

Named after the Milky Way, the fastest computer in the world is the National Supercomputer Center's Tianhe-2 in Guangzhou, in China. With peaks of 33.8 PFlops / s, it is powered by 3.12 million Intel Xeon E5-2692v2 processor cores. It is used for national security, and plans to double its size in the forthcoming years.

N^o2 **TITAN OAK RIDGE NATION LAB** UNITED STATES



A Cray XK7 system at the US Department of the Oak Ridge National Observatory in Oak Ridge, Tennessee, it uses Nvidia Tesla GPUs and AMD

S

upercomputers are like good vintage wines: they follow one another but are very seldom alike, and only a handful of them reach the superstar status in the academic, industrial or fundamental research fields. The ten supercomputers presented in these pages are the best elite machines and top

performers created by man. These supercomputers are also the most inventive in design and architecture. Because the slightest millisecond in data transfer or treatment or clock frequency can represent several million dollars in this breakneck speed race to performance. This ultralimited membership club only accepts the best of the best.

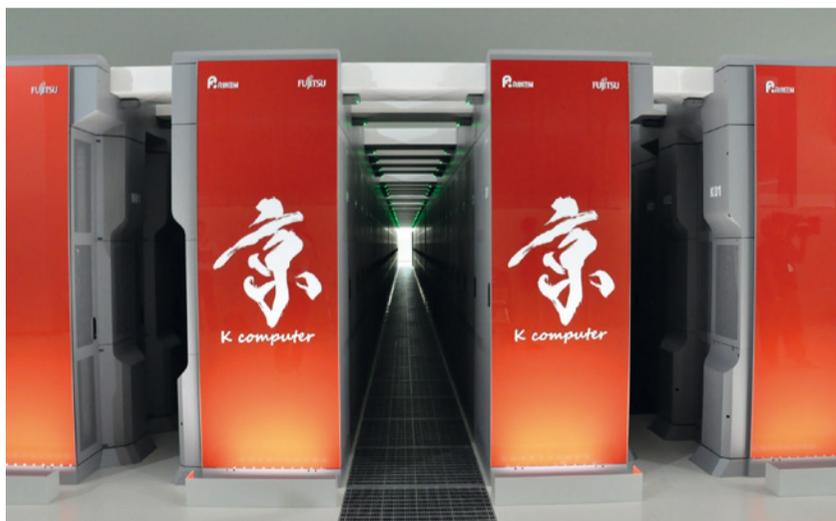
Opteron CPUs to create a 17.5 petaflop system for a range of science projects. Titan will be replaced in 2018 by IBM's Summit.

N°3 SEQUOIA LAWRENCE LIVERMORE NATIONAL LABORATORY UNITED STATES



Sequoia is another Blue Gene / Q system. Sequoia is installed at the Lawrence Livermore National Laboratory in California in the United States. It has the second biggest count of cores on the list, using 1 572 864 of them to produce 17 173 teraflops.

N°4 K COMPUTER FUJITSU JAPAN



The longest continuous presence in the Top 10 is Fujitsu K Computer, which achieves 10 510 teraflops of performance, despite being one of the less energy efficient systems on the list.

Japan's highest-ranking HPC, Fujitsu's K computer sits in Kobe, Japan's RIKEN Advanced Institute for Computational Science, where the machine is used to solve energy, sustainability, healthcare, climate change, industrial and space challenges.



N°5 MIRA ARGONNE LEADERSHIP COMPUTING UNITED STATES

Mira is one of the two Blue Gene / Q systems in this new top 10, boasting 8 596 teraflops. Mira is managed by the US Department of Energy laboratory located in Argonne suburban Chicago.



N°6 TRINITY LOS ALAMOS NATIONAL LABORATORY UNITED STATES

Trinity is the second of two newcomers in the top 10, taking the sixth place with a power of

8100 teraflops. It is the new main supercomputer in the Los Alamos, New Mexico National Laboratory.



N^o7 **PIZ DAINT** **SWISS NATIONAL SUPERCOMPUTING CENTER** SWITZERLAND

The alpine supercomputer named Piz Daint operates at 6271 teraflops, partly thanks to Nvidia Tesla GPUs with 73.808 active cores as accelerators. It is the most energy efficient system in this top 10, delivering 2697.2 megaflops per watt of power.



N^o8 **HAZEL HEN** **UNIVERSITY OF STUTTGART** GERMANY

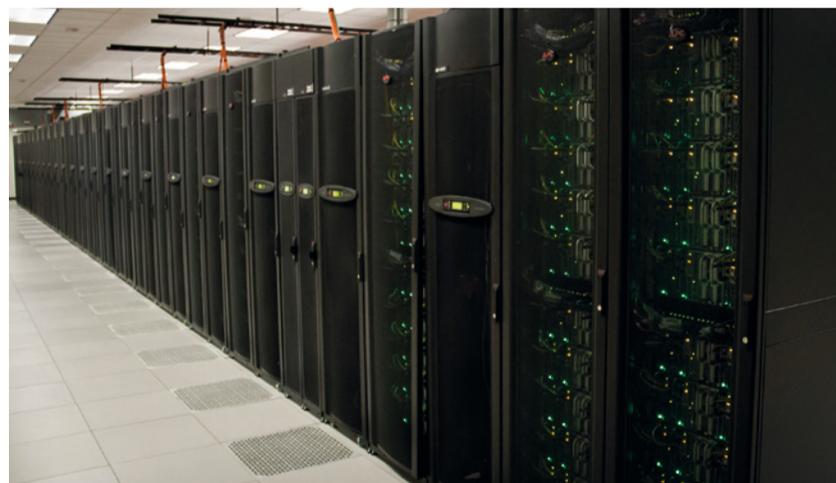
The first of the new entrants, Hazel Hen is the supercomputer at the University of HPC Stuttgart. Its 5 640 teraflops help Germany retain a place in the top 10.



N^o9 **SHAHEEN II** **KING ABDALLAH UNIVERSITY** SAUDI ARABIA

Shaheen II, an academic supercomputer located at the King Abdullah University in Saudi Arabia of Science and Technology, made the top 10 for the second consecutive time. It features 196 608 Intel Xeon E5 processor cores with 5.537 teraflops.

N^o10 **STAMPEDE** **UNIVERSITY OF AUSTIN** **TEXAS** UNITED STATES



The Stampede supercomputer of the University of Texas dropped to 10th place on this year's list, but it still remains a reference for researchers of this institution, with 5 186 teraflops of processing power.



LTO

7th GENERATION IS HERE!

According to the Linear Tape Open (LTO) consortium, the death of tape backup has been greatly exaggerated. Even in these times where storage seems to blend flash and Cloud based storage, tape still has a bright future ahead of it as it fills needs unmet by other media. Explanations.





“LTO-7 MORE THAN DOUBLES THE CAPACITY OF THE TAPE CARTRIDGE FROM THE PREVIOUS GENERATION, INCLUDING CAPACITIES UP TO 15 TB COMPRESSED PER CARTRIDGE.”

Three representatives of the LTO consortium agreed to answer our questions about the future of tape storage, on the eve of the availability of the seventh generation of the LTO format that will celebrate its fifteen years. They also offer us a glimpse of the next three generations. It is therefore the joint responses of Laura Loredo, Worldwide Director of Enterprise Storage and Servers at Hewlett Packard Enterprise LTO consortium representative, Shawn Haze, Business Line Manager, Data Protection & Retention at IBM and Terry Cochran, responsible Marketing and Communications at Quantum Corporation who answered our questions.

What are the main characteristics of LTO-7 compared to previous generations?

LTO-7 more than doubles the tape cartridge capacity from the previous generation, including capacities of up to 15TB* per cartridge when compressed. The specifications include a doubling of read/write heads in an advanced servo format to help achieve higher track density, which means that more data may be written to the same amount of tape within the cartridge. New formulation advancements also provided stronger magnetic properties, helping to increase overall capacity.



Laura Loredo is Worldwide Director of Enterprise Storage and Servers at Hewlett Packard Enterprise and the LTO Consortium representative.

Large files will also transfer more quickly on LTO-7 technology, with tape drive data transfer rates that are up to 750MB* per second, translated to more than 2.7 terabytes of data an hour per drive.

How does LTO-7 fit into storage architecture nowadays?

This next generation of LTO Ultrium technology, like those before it, works either as a standalone storage solution, or alongside cloud or disc in a hybrid environment. We're finding that many users look to tape for storing and archiving huge files and amounts of content (such as 4K video) that don't need to be accessed regularly.

A great example of how tape is being used is by cloud providers. Many cloud-based com-



“LTO TECHNOLOGY IS KNOWN AS AN ACTIVE ARCHIVE SOLUTION - MEANING IT WORKS BOTH AS COLD STORAGE/ARCHIVE, AS WELL AS AN ACTIVE, EVERYDAY STORAGE SOLUTION.”

panies – such as a popular email provider/host – use LTO Ultrium format tape drives to back up the data hosted on the cloud. So when the cloud servers go down, data is not lost but instead, easily found and restored.

Is LTO-7 mostly adapted for near-line or long term storage, and why?

LTO technology is known as an active archive solution - meaning it works both as cold storage/archive, as well as an active, everyday storage solution. LTO Ultrium format tape drives now have the Linear Tape File System (LTFS), which revolutionized the way tape is used in a near-line capacity. So while tape is known for being easily stored offline and off-site to guard against online data corruption and natural disasters, it can also be accessed for regular file recovery, just as easily as disk storage.

Can you describe some typical use cases in which LTO tape technology is best suited for?

Linear Tape Open Technology is a powerful choice for a variety of environments. The technology is the backbone of a generation of super tape storage products that feature high capacity and performance. LTO technology is used in industries including media and entertainment, law enforcement, government and security, HPC, and more.

In addition to the cloud server backup mentioned above, we're also seeing many law enforcement agencies turn to low cost tape technology to store the huge amount of video that



**Shawn Haze is
Business Line
Manager, Data
Protection &
Retention at IBM**

is streamed on police officer body cameras. Huge amounts of data that needs to be stored for compliance and potential access down the line, and done at a low, less-than-\$.01-per gigabyte cost.

Is LTO-7 still relevant in modern high-speed usage scenarios?

I am thinking about Big Data, Hadoop, Machine Learning for example. LTO is very relevant, if not critical, for successful storage in modern, high-speed usage scenarios. Tape is the most cost-effective media for storing mass quantities of data, especially in HPC and big data environments. When using peak programs and equipment such as Hadoop, it's important to demand peak performance from your data. When the value of data is at an all time high, it is imperative to have high performing, reliable storage. In any modern storage

scenario, tape provides a low cost tier of storage that is easily managed within the unstructured file system. In a traditional, multi-copy storage architecture, lowering the cost to store the data while still allowing near instantaneous access is critical, tape is the best choice for second and or third copy of critical data in analytic workflows.

What interfaces are available for LTO-7 units?

LTO-7 tape drives are offered in several interfaces. Stand alone units are offered in Half high for factor with Serial Attached SCSI (SAS) 6Gb connectivity. SAS 6Gb is compatible with SAS3Gb, 6 Gb and 12Gb host bus adapters (HBA). Automation based LTO-7 solutions offer both Full high and Half high form factors available in both Fibre Channel (FC) 8Gb and SAS 6Gb. The flexibility of the LTO-7 tape technology has enabled third party manufacturers to offer other interfaces for connectivity in a variety of manners.

When will the first LTO-7 media be available?

LTO-7 media has been licensed for production by Sony and Fujifilm after passing interchange testing, so units will be available shortly, though that is dependent on the timelines of each individual manufacturer.

Previous generations are still available from any of the major manufacturers, and we encourage users to only purchase the technology with the Ultrium format trademark. The Ultrium format trademark identifies products that have passed annual third-party compliance verification to ensure compliance with the Ultrium tape format specification. Use of this trademark is granted only after an Ultrium tape format product has passed the



Terry Cochran is Marketing Communications Manager at Quantum Corporation

compliance verification testing process.

Who are the manufacturers for LTO products?

The following companies are current LTO Format Specification Participants that have delivered LTO compatible products and components: Atmel, Fujifilm, Hewlett Packard Enterprise, IBM, Imation, Maxell, Mountain Engineering, Netc, NXP Semiconductors, Quantum, Sony, Spectra Logic, TDK, Teijin Dupont Film and Toray

Industries. Fujifilm and Sony have both passed compliance verification requirements to sell LTO-7 cartridges.

In terms of roadmap, what lies ahead for LTO tape technology? (capacity and evolutionary steps and dates?)

The LTO Ultrium roadmap addresses complex storage needs over time, and builds out a plan to continually increase capacity and speed. LTO generation 7 technology offers an unprecedented 15TB compressed capacity and transfer speeds up to 750MB/s. While the LTO Program is thrilled to offer such a high degree of capacity and speed, we also look to the future. LTO generation 8 is slated to offer up to 32TB compressed capacity and up to 1180MB/s transfer speed. LTO generation 9 goes further with up to 62.5TB capacity and up to 1770MB/s transfer, while LTO generation 10 will offer up to 120TB capacity and up to 2750MB/s. Each new generation of LTO technology nearly doubles capacity and speed capabilities, and new generations are released roughly every two years.



NEWSFEED



HPC
REVIEW



MIRA

THE WORLD'S FIRST
LHC SIMULATOR
SUPERCOMPUTER

CONNECTING THOUSANDS OF COMPUTERS AND STORAGE SYSTEMS ACROSS 41 COUNTRIES, THIS INTERNATIONAL GRID COMPUTING INFRASTRUCTURE ALLOWS DATA TO BE VIEWED AND ANALYZED IN NEAR REAL TIME BY AN INTERNATIONAL COMMUNITY OF MORE THAN 8,000 PHYSICISTS.

Physicists at Argonne Computing Center use their Mira Supercomputer to perform experiment simulations of the Large Hadron Collider (LHC) to interpret future LHC data. Researchers at Argonne Leadership Computing installation (ALCF) helped the team to optimize their code for the supercomputer, which allowed them to simulate billions of particle collisions faster than ever.

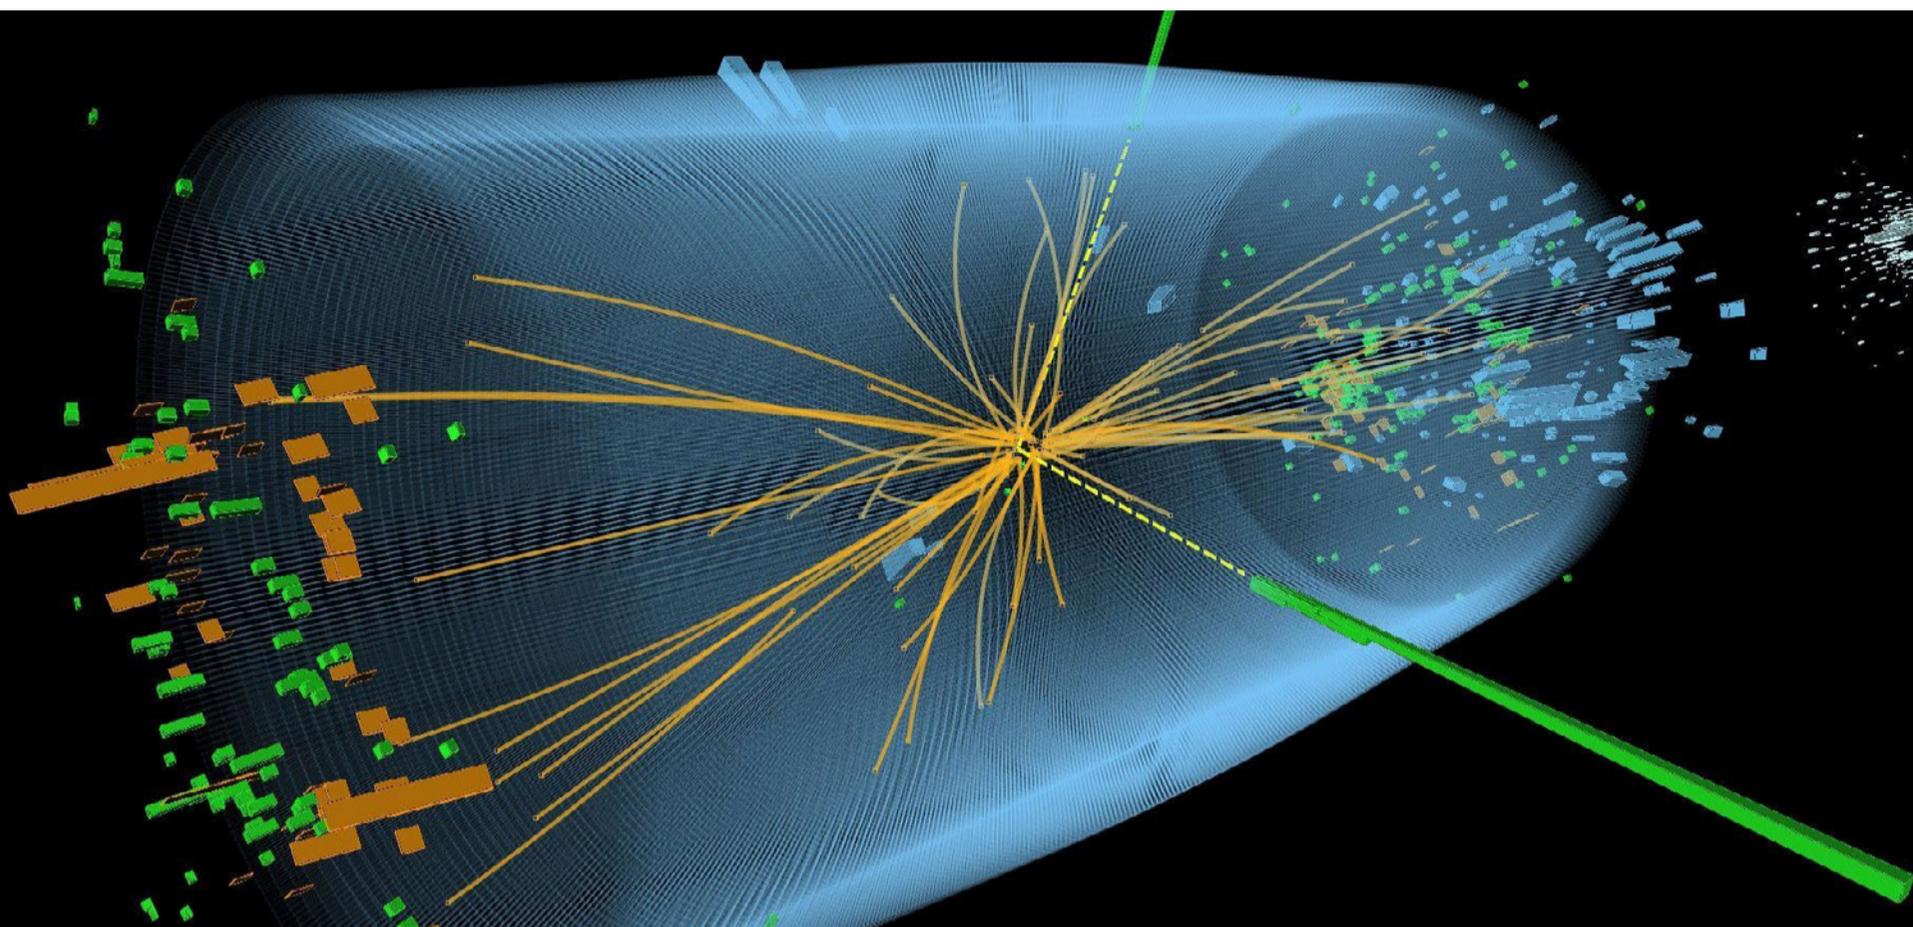
BILLIONS OF PARTICLE COLLISIONS PER SECOND ANALYZED

The Large Hadron Collider at CERN (LHC) is the world's most powerful particle accelerator, scientists launch billions of particle collisions per second in their quest to understand the fundamental structure of matter. With each single particle collision producing about one megabyte of data, the LHC's facility, located on the border of France and Switzerland, generates a massive amount of data. Even after filtering approximately 99% of the data, scientists have annually about 30 petabytes (or 30 million gigabytes) of data to analyze in a wide range of physics experiments, such as studies on the Higgs Boson and the Black Matter. To meet the considerable challenge of interpreting this amount of data, researchers at the US Department of Argonne National Laboratory

have demonstrated the simulation of particle collisions with Mira, the 10 petaFLOP IBM Blue Gene / Q supercomputer. "Simulating collisions is essential to help us understand the response of particle detectors," said lead researcher Tom LeCompte, Argonne physicist and former coordinator for the ATLAS experiment of the LHC, one of the four particle detectors. "The differences between the simulated data and experimental data can lead us to discover signs of new physics."

WORLD PREMIERE

This successful simulation marks the first time a supercomputer was used successfully to perform massively parallel simulations of the LHC particle collisions. This demonstrates that supercomputers can help drive future discoveries at the LHC while accelerating the pace at which the simulated data can be produced. The project also shows how IT resources can be used to facilitate the culmination of physics experiments. Since 2002, scientists at the LHC have relied on the Worldwide LHC Computing Grid for all their needs of data processing and simulation. Connecting thousands of computers and storage systems across 41 countries, this international grid computing infrastructure allows data to be viewed and analyzed in near real time by an international community of more than 8,000 physicists working on the four major LHC experiments. "Grid computing has been a success for the LHC, but there are still limits," LeCompte said.



“The first is that some of LHC events simulations are so complex that it would take weeks to complete. The second is that the LHC computing needs require to be increased tenfold in the coming years.” To investigate the use of supercomputers as a possible tool for the LHC, LeCompte sought and obtained computing time from the ALCF Advanced Computing Research Centre. His project focuses on the simulation of ATLAS events that are difficult to simulate with CERN’s existing grid computing infrastructure. Although the volume and nature of data from the LHC seem a natural choice for one of the fastest supercomputers in the world, it took extensive work to adapt an existing LHC simulation method for Mira’s massively parallel architecture. With the help of ALCF researchers Tom Uram, Hal Finkel and Venkat Vishwanath, the Argonne team has transformed ALPGEN, an application based on the Monte Carlo algorithm that generates the events in hadron collisions, from a single threaded simulation code into a massively multi-threaded code that could function effectively on Mira. By improving I/O performance of the code and reducing its memory use, they have managed to adapt ALPGEN to

operate on the Mira system and run code 23 times faster than the original. This code optimization work allowed the team to simulate millions of LHC collision events in parallel. “In performing these tasks on Mira, the team conducted the equivalent of two years of ALPGEN simulations in weeks, and the CERN’s LHC Grid Computing Infrastructure became available to perform other tasks,” said Uram.

If supercomputers like Mira are better integrated into the workflow of the LHC, LeCompte hopes that many more simulations could later be transferred to supercomputers, relieving the LHC computing infrastructure and accelerating significantly the simulations with the collected data. To move in that direction, his team plans to increase the range of executable code on Mira, the next candidates being Sherpa, another event generation code, and Geant4, a code to simulate the passage of particles through matter. “We also plan to help other physics groups to use supercomputers as Mira” said LeCompte. “According to our own experience it takes a year or two to adapt and rewrite the original code, and another year to run it on a large scale.”



NVIDIA OPENS THE WAY OF HYPERSCALE MACHINE LEARNING

N

VIDIA announced the Tesla M4 and M40, two cards dedicated to accelerate intensive Machine Learning workloads. The M40 is aimed at researchers to help them create new Neural Networks for each application they want to operate with artificial intelligence. The M4 is a low power accelerator designed to deploy these neu-

ral networks across the data center. Both cards benefit from a specific library for accelerating GPU treatments. Together, they enable developers to use the Tesla accelerated computing platform to manage machine learning on the data center. “The race to artificial intelligence is launched” said Jen-Hsun Huang, co-founder and CEO of NVIDIA. “Machine Learning is undoubtedly one of the most important developments at present in the IT field, the same as the computer, Internet and the Cloud. Machine learning is the big IT challenge of our generation. We created the range of Tesla HyperScale acce-

lerators to multiply by 10 the Machine Learning algorithms calculation power. Time and money earned for datacenters are huge” he says.

ACCELERATING WEB APPLICATIONS

These new software and hardware products are specifically designed to accelerate the flow of Web applications that integrate artificial intelligence capabilities. Significant advances allow to use artificial intelligence techniques to create smarter services and applications. Machine learning is used to achieve accurate speech recognition. It also allows to create automatic recognition algorithms in videos or photos, with possibility for future research. Automatic learning also enables facial recognition in videos or pictures even when the face is veiled in part. Finally, it operates services that help to know the tastes and personal interests of the user, arranging schedules, provide relevant information and answer precisely a voice command. The challenge is to get the awesome power IT needed to innovate and to train neural networks and to process information to respond immediately to the billions of requests from consumers.

TESLA M40: DESIGNED FOR DATA SCIENTISTS

With 3072 stream processors and 7 TFlops of computing power, the NVIDIA Tesla M40 accelerator GPU allows data scientists to save time, days or weeks, when they train their deep neural networks to manage huge amounts data while obtaining a higher accuracy. Adapted to automatic learning, it divides according to the manufacturer, the training time compared with 8 CPU processors (1.2 days against 10 days using AlexNet for a classical training). The support for NVIDIA GPUDirect enables rapid training and creation of multinode neural networks.

TESLA M4 FOR THE DATACENTER

The M4 NVIDIA Tesla GPU is a low power accelerator created for HyperScale environments to enable responsive service and demanding web applications, as video transcoding, processing images and videos, and automatic

learning inference. According to Nvidia, the M4 is capable of improving data analysis up to 5 times more video streams simultaneously than a conventional CPU. On the energy consumption side, the Tesla M4 consumes 50 to 75 watts, and provides energy efficiency up to 10 times better than a CPU while processing video and machine learning algorithms. The M4 has 1024 stream processors and 2.2 TFlops calculation power. The onboard memory is 4GB against 12GB on the M40.

NVIDIA HYPERSCALE SUITE, A DEDICATED SOFTWARE KIT

The new software suite contains NVIDIA HyperScale tools for both developers and data center managers and is specially designed for deploying web services, such as:

- **CUDNN** - the most popular software algorithms that manages deep convolutional neural networks used for artificial intelligence applications.
- **ACCELERATED FFmpeg MULTIMEDIA SOFTWARE GPU** - accelerated transcoding and processing videos.
- **NVIDIA GPU ENGINE REST** – Create easily accelerated and deploy web services high processing capacity and low latency ranging from dynamic resizing Image accelerating research through the image classification.
- **NVIDIA IMAGE RESIZING ENGINE** - GPU-accelerated APIs with REST Service which allows to resize images 5 times faster than a CPU.

MESOSPHERE SUPPORT

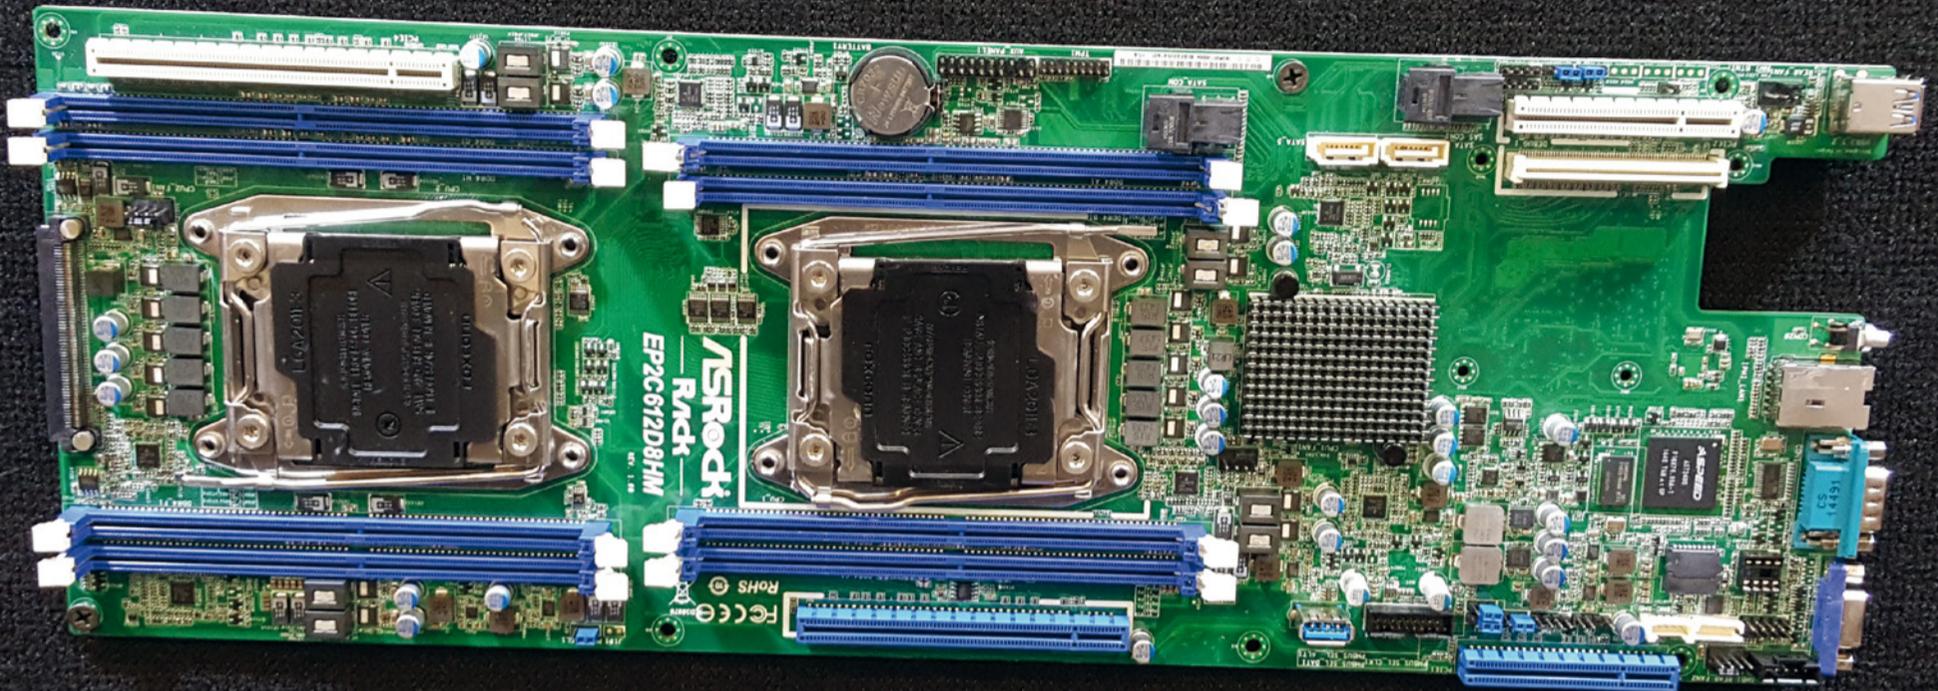
At the last demonstration of Tesla accelerated computing platform NVIDIA announced its collaboration with Apache to add GPU technology support within the Mesosphere Datacenter Operating System (DCOS). This strategic partnership will allow Web service companies to build and deploy their new generation accelerated applications more easily in data centers. The Tesla M40 Accelerator and HyperScale Suite are available, whereas the Tesla M4 Accelerator will be available in the first quarter of 2016.



ASROCK RACK STRATEGY

**CUSTOMIZED AND GREEN
HPC SERVERS**

WEISHI SA, ASROCK RACK GENERAL MANAGER
INTERVIEW BY **FRÉDÉRIC MILLIOT**



PC Review has interviewed Mr. Weishi Sa, ASRock Rack Worldwide General Manager who updated us on ASRock Rack product portfolio and strategy, including some geographical and distribution figures and even a glimpse into ASRock Rack's Research and Development strategy.

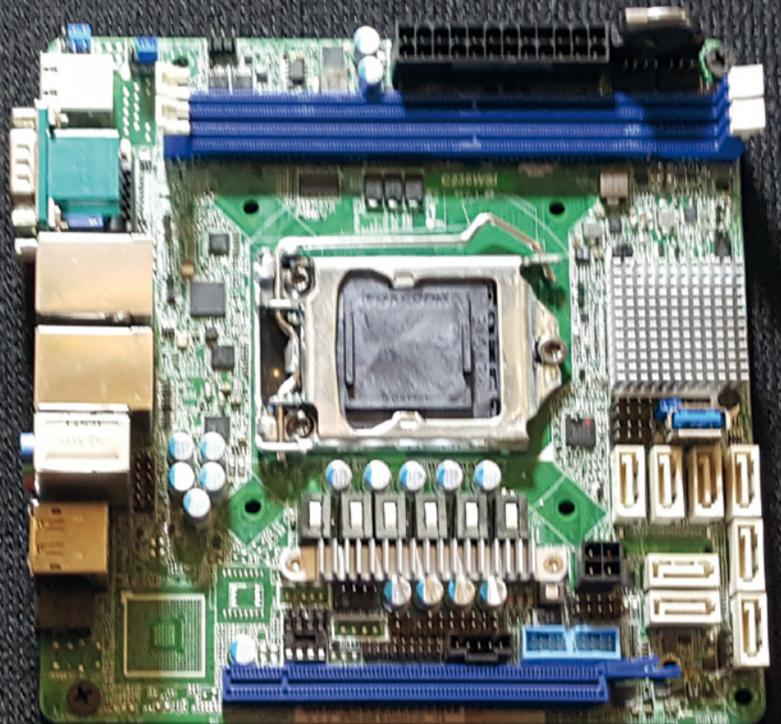
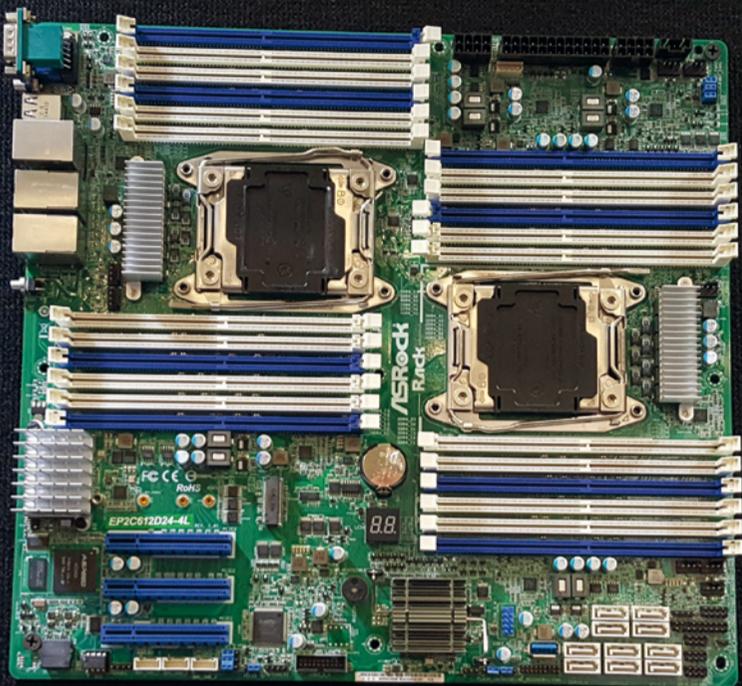


Can you please explain us your Xeon D Server motherboard launch and the planned availability of your Xeon D servers ?

As an innovative server technology provider, we focus on high performance and high-efficiency server products such as HPC, Cloud Computing and Datacenters. Our latest D1540D4X

motherboard series support Intel Xeon D processors today. The new series completely utilizes Intel's latest processor technology to offer perfectly optimized computing solutions.

Intel Xeon D product families, as the first Intel Xeon SoC, pushes the performance, efficiency and the TCO to the next level with greater density and integration. The upgraded



14nm processor technology and higher levels of I/O integrations scales the performance and reliability to lower power design points, yet includes the datacenter processor features. ASRock Rack's latest D1540D4X series features strong computing performance in E5 level, memory capacity and satisfying I/O expandability. Our Xeon D Servers are available now.

What are the market segment / use cases you are targeting with your Xeon D motherboard Series ?

This series is ideal for HPC-based micro servers based on its form factor. It also gives the benefits for users to build up substantial computing nodes under a cloud environment. Besides, the series is perfect for hyperscale storage servers, thanks to SoC integration. Moreover, this series is a comprehensive solution for mid-communication servers under a

cost-effective yet reliable condition. Overall, D1540D4X takes the full advantage of Intel's latest processor technology, offering a low power and compact server solution for customers to cope with the existing critical issues of Datacenters and Cloud environment.

Can you update us on your server strategy worldwide ?

We sell to wholesale distributors like Super-Micro. We also do OEM mainly in China and ODM in Russia We also have a local distributor ecosystem for Europe and France.

What is your geographical revenue breakdown ?

The revenue percentage is about 30% in Europe, 10% in Russia, 20% in Japan / Asia and the rest in the United States. That helps us shorten the shipping times across our entire product range.



THESE BOARDS ARE IDEAL FOR NETWORK AND WARM STORAGE DEVICES AS WELL AS MICRO-SERVERS. GIVEN THE LOW POWER NATURE OF THE XEON D PLATFORM, SOME OF THEM CAN ALSO BE USEFUL IN HOME LAB SETTINGS FOR EXPERIMENTING WITH VIRTUALIZATION OR EVEN ACT AS BOARDS FOR HIGH-END DEVELOPMENT MACHINES.

Can you introduce us your greenest server?

HPC servers usually consume much power. However, energy efficiency is at the core of our business. Our customers are increasingly looking at the energy / performance efficiency ratio.

Can you let us know your customers top use cases ?

We work very closely with the academic sector. University computation centers are among the top users of our server products. Customers want integrated server packages, mostly in Europe and Japan. Whereas in Asia and Japan, they have a more customized approach.

What is your future product strategy ?

We have a strategy and product roadmap up to 3 years. We are confident in that HPC will continue growing in demand. In order to increase our server versatility and usefulness, some customers in China often add GPU accelerators in our servers for their needs, like NVIDIA Tesla K40 GPUs.

How much do you invest in Research and Development ?

We invest 30% of our global revenue into HPC Research and Development. We certify our high density servers with NVIDIA Tesla GPU accelerators. ASRock Rack's Xeon D lineup consists of one board using the Xeon D-1520 and six boards using the Xeon D-1540. Customers have the option of going with either the mini-ITX (mITX) form factor or the micro-ATX (uATX) form factor. The mITX boards are all compatible with 1U rackmount chassis.

In addition to the motherboard size, the differentiation aspects come in the form of support for different varieties of LAN ports, PCIe slot configurations, additional storage ports using the LSI 3008 HBA and different USB 3.0 port configurations. Unlike the mITX boards, all the uATX boards come with a COM port in the rear I/O. The following tables summarize the features of the various products in the ASRock Rack Xeon D lineup.

These boards are ideal for network and warm storage devices as well as micro-servers. Given the low power nature of the Xeon D platform, some of them can also be useful in home lab settings for experimenting with virtualization or even act as boards for high-end development machines.



XEON D PRODUCT PORTFOLIO

ASRock Rack's Xeon D lineup consists of one board using the Xeon D-1520 and six boards using the Xeon D-1540. Customers have the option of going with either the mini-ITX (mITX) form factor or the micro-ATX (uATX) form factor. The mITX boards are all compatible with 1U rackmount chassis. In addition to the motherboard size, the

differentiation aspects come in the form of support for different varieties of LAN ports, PCIe slot configurations, additional storage ports using the LSI 3008 HBA and different USB 3.0 port configurations. Unlike the mITX boards, all the uATX boards come with a COM port in the rear I/O. The following tables summarize the features of the

various products in the ASRock Rack Xeon D lineup. These boards are ideal for network and warm storage devices as well as micro-servers. Given the low power nature of the Xeon D platform, some of them can also be useful in home lab settings for experimenting with virtualization or even act as boards for high-end development machines.

mITX BOARDS	D1520D4I	D1540D4I	D1540D4I 2L2T
SIP	Intel Xeon D-1520	Intel Xeon D-1540	
RAM	4x DDR4 DIMM Slots 2133 / 1866 MHz RDIMMs (Up to 128 GB)		
PCIe EXPANSION SLOTS	1x PCIe 3.0 x16		
STORAGE CONTROLLERS	6x SATAIII 6 Gbps from integrated PCH in the Xeon D SiP (4x via mini-SAS connector) (1x with SATA DOM support) 1x SATAIII 6 Gbps from Marvell 9172 (via M.2 2280 interface)		
LAN CONTROLLERS	2x RJ45 1GbE (Intel i210)		2x RJ45 1GbE (Intel i210) 2x RJ45 10GbE (Intel X557-AT2)
BOARD MANAGEMENT CONTROLLER	ASPEED AST2400		
IPMI LAN CONTROLLER	1x Realtek RTL8211E		
DISPLAY OUTPUT	1x D-Sub VGA		
USB PORTS	2x USB 3.0 Type-A (Rear I/O)		

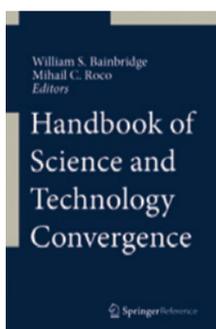
uATX BOARDS	D1540D4U-2T8R	D1540D4U-208R	D1540D4U-2T208R	D1540D4U-2L+
SIP	Intel Xeon D-1540			
RAM	4x DDR4 DIMM Slots 2133 / 1866 MHz RDIMMs (Up to 128 GB)			
PCIe EXPANSION SLOTS	1x PCIe 3.0 x8 (x16 physical) 1x PCIe 3.0 x8 (x8 physical)		1x PCIe 3.0 x16 1x PCIe 3.0 x8	
STORAGE CONTROLLERS	6x SATAIII 6 Gbps from integrated PCH in the Xeon D SiP (4x via mini-SAS connector) (1x with SATA DOM support) 8x SAS3 12Gbps from LSI 3008 HBA (via mini-SAS HD connector) 1x SATAIII 6 Gbps from Marvell 9172 (via M.2 2280 interface)			
LAN CONTROLLERS	2x RJ45 10GbE (Intel X550)	2x 10G SFP+ Fiber	2x 10G SFP+ Fiber 2x RJ45 10GbE (Intel X540)	2x RJ45 1GbE (Intel i350)
BOARD MANAGEMENT CONTROLLER	ASPEED AST2400			
IPMI LAN CONTROLLER	1x Realtek RTL8211E			
DISPLAY OUTPUT	1x D-Sub VGA			
USB PORTS	2x USB 3.0 Type-A (Rear I/O) 1x USB 3.0 Type-A (Internal Connector) 1x USB 3.0 Header			



books

HANDBOOK OF SCIENCE AND TECHNOLOGY CONVERGENCE

William Bainbridge, Mihail C. Roco
1180 pages, Springer International Publishing
US \$599



Scientists and engineers have long been aware of the tension between narrow specialization and multidisciplinary cooperation, but now a major transformation is in process that will require technical fields to combine far more effectively than formerly in the service of human benefit. This handbook will catalog all the ways this can be accomplished and the reasons it must be. Nature is a single coherent system and diverse methods of scientific and engineering investigations should reflect this interlinked and dynamic unity. Accordingly, general concepts and ideas should be developed systematically in in-

terdependence, with cause-and-effect pathways, for improved outcomes in knowledge, technology and applications. At the same time, industrial and social applications rely on integration of disciplines and unification of knowledge. Thus, convergence is both a fundamental principle of nature and a timely opportunity for human progress. This handbook will represent the culmination of fifteen years of workshops, conferences and publications that initially explored the connections between nanotechnology, biotechnology, information technology and new technologies based on cognitive science. A constant emphasis on human benefit then drew in the social sciences, even as shared scientific and ethical principles brought in sustainability of the Earth environment and the challenge of equitable economic advancement. The intellectual contributions of literally hundreds of scientists and engineers established a number of research methods and analytical principles that



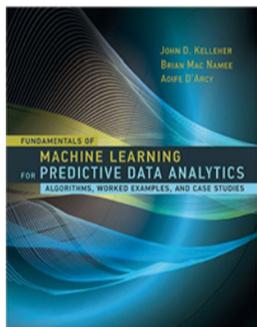
could unite disparate fields. The culmination has been called Convergence of Knowledge and Technology for the benefit of Society (CKTS), defined as the escalating and transformative interactions among seemingly different disciplines, technologies, communities and domains of human activity to achieve mutual compatibility, synergism and integration.

FUNDAMENTALS OF MACHINE LEARNING FOR PREDICTIVE DATA ANALYTICS

John D. Kelleher, Brian Mac Namee, Aoife D'Arcy

MIT Press, 624 pages

US \$70 / eBook US\$ 70



Machine learning is often used to build predictive models by extracting patterns from large datasets. These models are used in predictive data analytics applications including price prediction, risk assessment, predicting customer behavior, and document classification. This introductory textbook offers a detailed and focused treatment of the most important machine learning approaches used in predictive data analytics, covering both theoretical concepts and practical applications. Technical and mathematical material is augmented with explanatory worked examples, and case studies illustrate the application of these models in the broader business context. After discussing the trajectory from data to insight to decision, the book describes four approaches to machine learning: information-based learning, similarity-based learning, probability-based learning, and error-based learning. Each of these approaches is introduced by a nontechnical explanation of the underlying concept, followed by mathematical models and algorithms illustrated by detailed worked examples. Finally, the book considers techniques for evaluating prediction models and offers two case studies that describe specific data analytics projects through each phase of development, from formulating

the business problem to implementation of the analytics solution. The book, informed by the authors' many years of teaching machine learning, and working on predictive data analytics projects, is suitable for use by undergraduates in computer science, engineering, mathematics, or statistics; by graduate students in disciplines with applications for predictive data analytics; and as a reference for professionals.

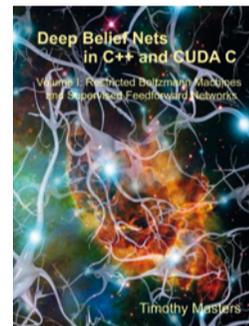
DEEP BELIEF NETS IN C++ AND CUDA C: VOLUME 1

RESTRICTED BOLTZMANN MACHINES AND SUPERVISED FEEDFORWARD NETWORKS

Timothy Masters

CreateSpace Independent Publishing Platform

244 pages, US \$45



Deep belief nets are one of the most exciting recent developments in artificial intelligence. The structure of these elegant models is much closer to that of human brains than traditional neural networks; they have a 'thought process' that is capable of learning abstract concepts built from simpler primitives. A typical deep belief net can learn to recognize complex patterns by optimizing millions of parameters, yet this model can still be resistant to overfitting. This book presents the essential building blocks of the most common forms of deep belief nets. At each step the text provides intuitive motivation, a summary of the most important equations relevant to the topic, and concludes with highly commented code for threaded computation on modern CPUs as well as massive parallel processing on computers with CUDA-capable video display cards. Source code for all routines presented in the book, and the DEEP program which implements these algorithms, are available for free download from the author's website. NOTE... The source code available for free download includes all of the code listed in the book, along with some libraries of related rou-



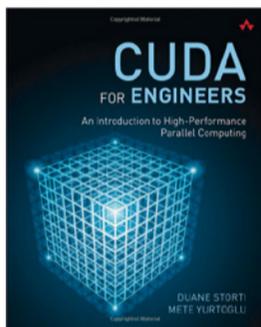
tines. Complete code for the DEEP program is not included; this code is enormous, as it includes many Windows-only interface routines, screen display code, and so forth. Users who wish to write their own DBN programs are responsible for implementing their own hardware/OS interface, while using the supplied code for the mathematical calculations.

CUDA FOR ENGINEERS AN INTRODUCTION TO HIGH- PERFORMANCE PARALLEL COMPUTING

Duane Storti, Mete Yurtoglu

Addison-Wesley Professional

352 pages, US \$45



CUDA for Engineers gives you direct, hands-on engagement with personal, high-performance parallel computing, enabling you to do computations on a gaming-level PC that would have required a super-computer just a few years ago. The authors introduce the essentials of CUDA C programming clearly and concisely, quickly guiding you from running sample programs to building your own code. Throughout, you'll learn from complete examples you can build, run, and modify, complemented by additional projects that deepen your understanding. All projects are fully developed, with detailed building instructions for all major platforms. Ideal for any scientist, engineer, or student with at least introductory programming experience, this guide assumes no specialized background in GPU-based or parallel computing. In an appendix, the authors also present a refresher on C programming for those who need it. Coverage includes : Preparing your computer to run CUDA programs / Understanding CUDA's parallelism model and C extensions / Transferring data between CPU and GPU / Managing timing, profiling, error handling, and debugging / Creating 2D grids / Interoperating with OpenGL to provide real-time user interactivity / Performing basic simulations with differential equations / Using sten-

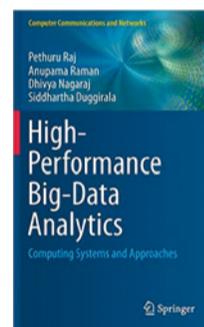
cils to manage related computations across threads / Exploiting CUDA's shared memory capability to enhance performance / Interacting with 3D data: slicing, volume rendering, and ray casting / Using CUDA libraries / Finding more CUDA resources and code. Realistic example applications include : Visualizing functions in 2D and 3D / Solving differential equations while changing initial or boundary conditions / Viewing/processing images or image stacks / Computing inner products and centroids / Solving systems of linear algebraic equations / Monte-Carlo computations

HIGH PERFORMANCE BIG DATA ANALYTICS COMPUTING SYSTEM AND APPROACHES

Pethuru Raj, Anupama Raman, Dhivya Nagaraj, Siddhartha Duggirala

428 pages, Springer International Publishing

US \$129 / eBook US \$99



This book presents a detailed review of high-performance computing infrastructures for next-generation big data and fast data analytics. Features: includes case studies and learning activities throughout the book and self-study exercises in every chapter; presents detailed case studies on social media analytics for intelligent businesses and on big data analytics (BDA) in the healthcare sector; describes the network infrastructure requirements for effective transfer of big data, and the storage infrastructure requirements of applications which generate big data; examines real-time analytics solutions; introduces in-database processing and in-memory analytics techniques for data mining; discusses the use of mainframes for handling real-time big data and the latest types of data management systems for BDA; provides information on the use of cluster, grid and cloud computing systems for BDA; reviews the peer-to-peer techniques and tools and the common information visualization techniques, used in BDA.

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.



INTERNET OF THINGS



**THE POWER
TO CHANGE
THE WORLD**

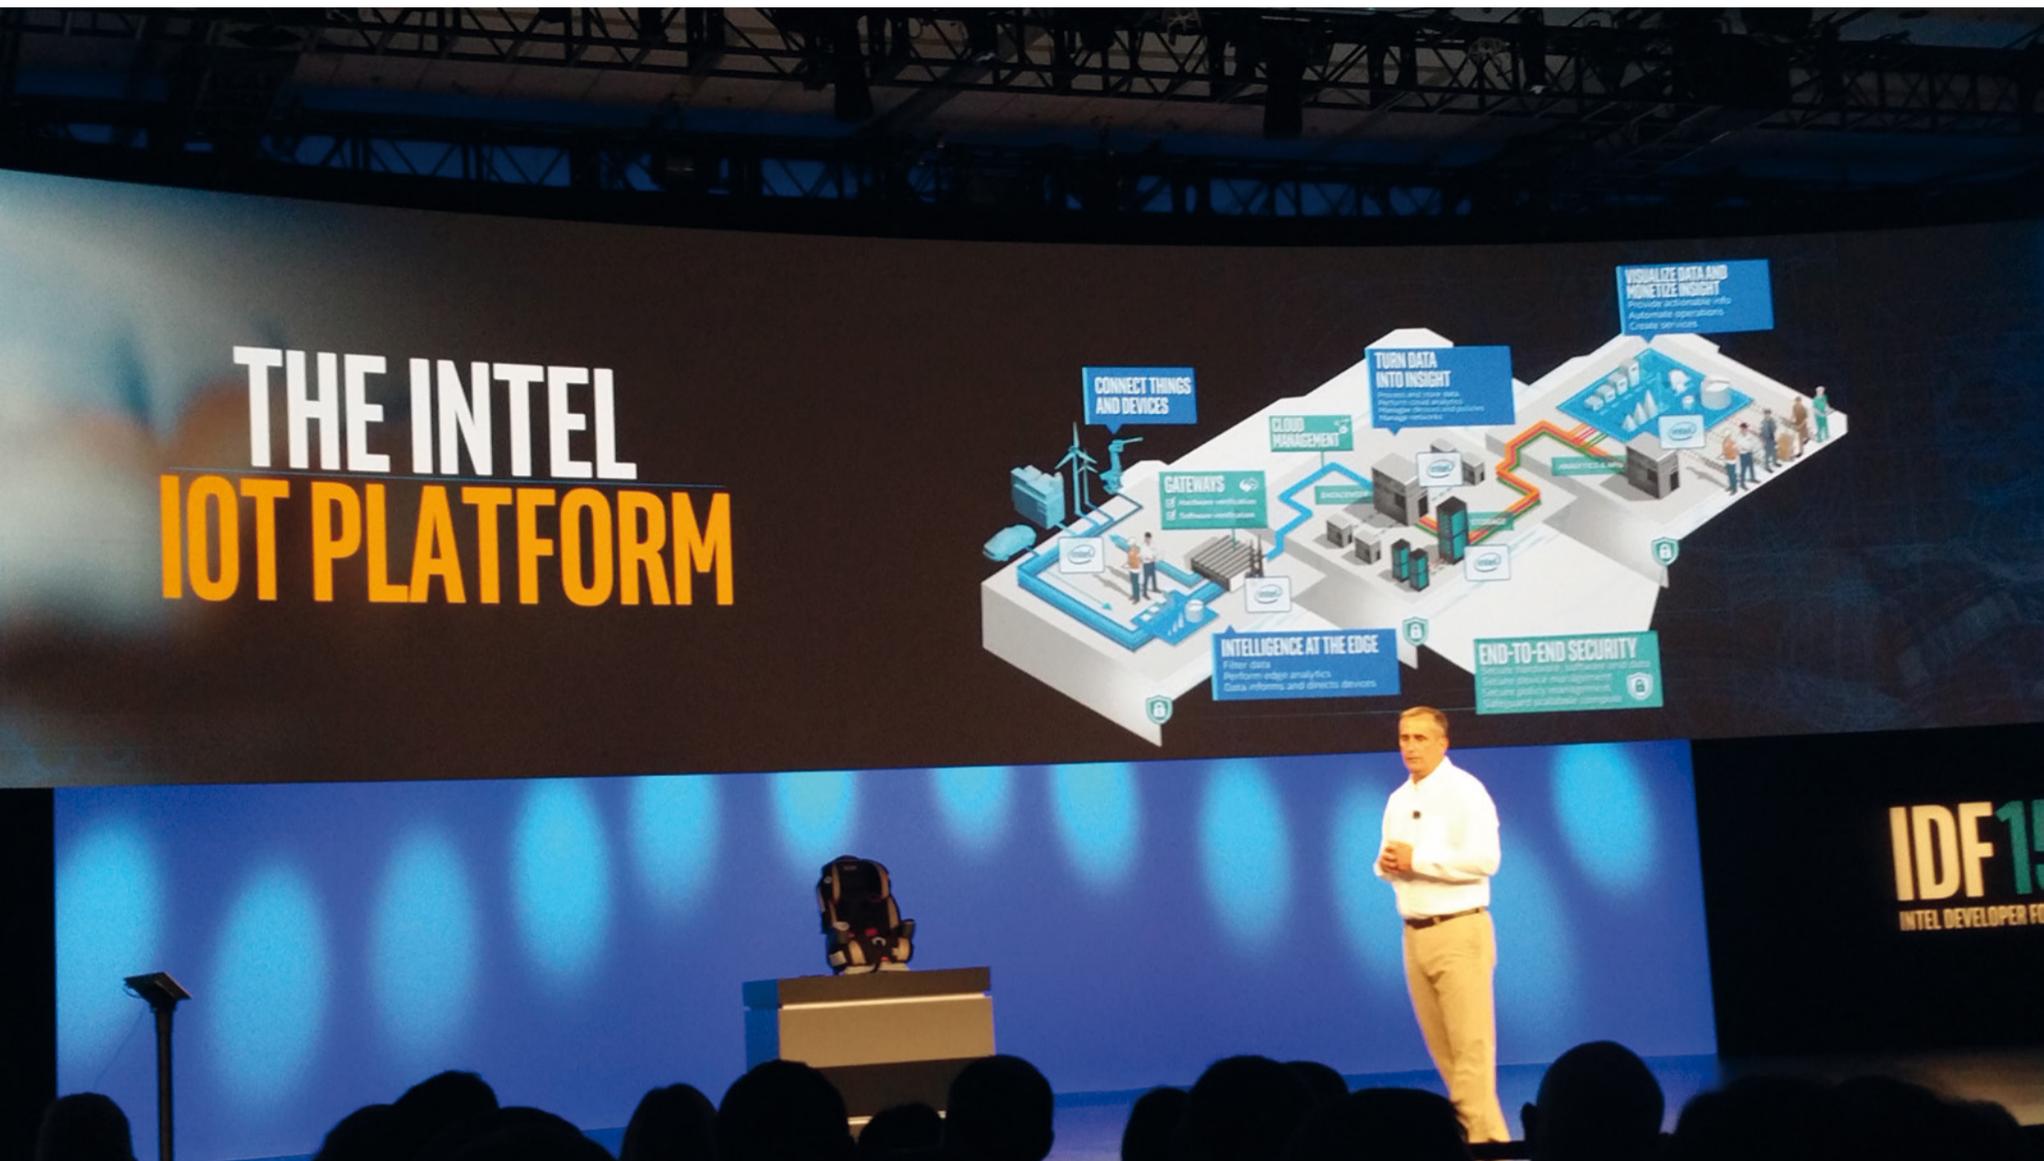


BIRTH OF AN
INDUSTRIAL
ECOSYSTEM



The Internet of Things was in dire need of a first rank actor to propose an overall industrial vision. It is to this position that Intel has stepped up to in the last months, with a 360° view that covers many aspects neglected or omitted by the protagonists of this emerging sector. The multiple actions and product launches these last years around IoT foreshadowed two possible outcomes. A top or a flop. With a tendency for the latter, as the abundance of creative but disorganized initiatives raised fears of incompatible standards and above all, riddled with risks and potential vulnerabilities after massive adoption which would mean a huge hangover for its users. A consensus had emerged: the IoT platforms were not safe!

Given the massive amount of personal data at risk, Intel reckoned it needed to act. Brian Krzanich, CEO of Intel, has therefore taken the lead in demonstrating several times since Intel Developer's Forum, an open initiative several floors and at the disposal of developers worldwide. Enough to ensure solid and reliable foundation to support the emergence The Intel Developers Forum is the place where you speak, see and experience the future of technology and innovation by looking at the past and future progress due to developments around Intel technologies "Intel builds Products and developers make them accessible to users," declared Krzanich during his opening address. Three days of preview of the new generations of Intel Products and a creativity at its highest level according to Intel's CEO "this is the first time that there are as many diversity in terms of opportunities for developers. We arrive at the age of customization compu-



ting...” And IoT participates significantly to this market development for cars, clothing, home...

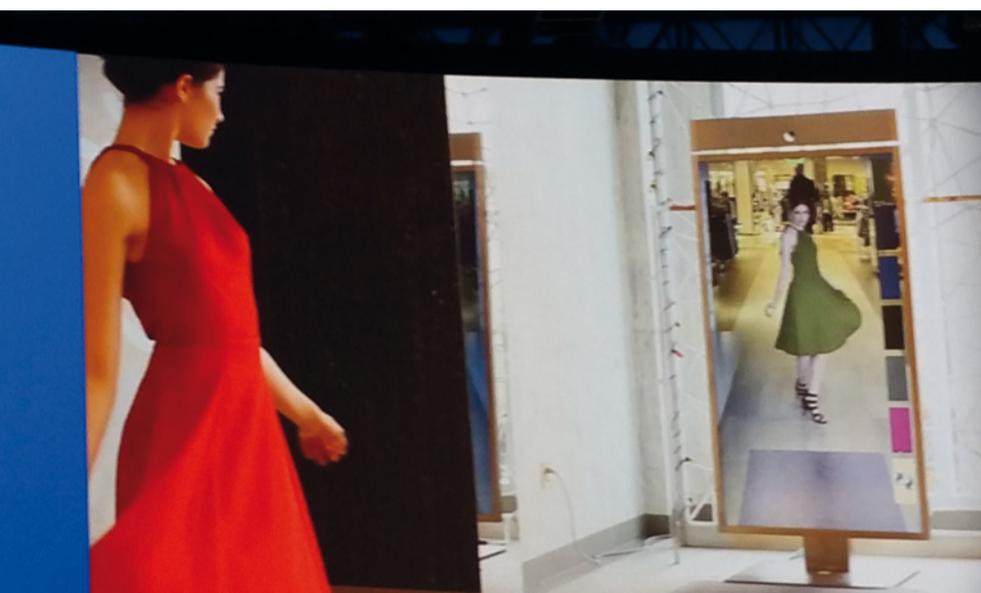
A POTENTIAL IN THREE AXIS

Intel categorizes this potential development in three axis that describe and guide innovation. The first is “Sensification” a neologism that designates the sound, the touch and the usual interface between user and machine. The need today is to have the same feelings between a machine and its user than that one would have with another human being. More natural sensations during exchanges with its IT equipment. Exchanges which are not necessarily at the initiative of the user, that do not require to push a button for the machine to respond to the interrogations of its user. The second axis is “Smart and Connected”. It refers to the “personalization” of the computing environment for a given user. Finally the

last assertion, “Extension of You” represents, according to Intel, the computer that becomes an extension of a person.

THE SENSIFICATION, SOME ACHIEVEMENTS

We talk and we hear: this basic principle gave birth to a product co-developed by Intel and Microsoft: Wake-on-Voice runs on Windows 10. The computing device remains listening permanently using the included audio technology in the processor and, in one low power state which does not exhaust the batteries. It works with Cortana, Microsoft’s digital personal assistant, competitor to Apple’s Siri, Google Now and Blackberry Wizard. Google, Intel tackled rather to the reduction of latency in the field of audio and it is here substantially closer exchanges audio in real time. Regarding Intel’s RealSense technology, it may allow, among other things, robots,



drones and generally any moving camera that incorporates this technology, avoid obstacles encountered on his path. Still using technology RealSense, Tango is a Google-Intel co-developed project which opens new possibilities for developing applications on smartphones. For example, in the field of augmented reality, scanning physical environments in 3D using a smartphone almost in real-time. At the heart of all these innovations, the equipment like the software are often open source. And in the case of the Jeeves home robots, Robotics Foundation made the Opensource part of the endeavour. In connection with the RealSense technology, a new ecosystem has formed around Intel with over 13 supported platforms: Linux, Android, Ross, Windows 10, Unity, Twitch (company that offers the technology to stream video games, VoD and social gaming), Scratch, OSVR, Xsplit, Structure SDK, project tango, Unreal engine ... Sensification is also prevalent with the focus on sensors with respect to video games, physical and virtual. The idea is to change the relationship between the gamer and play. Everything is in the felt as demonstrated by the game co-developed by iRacing and VRX. Beyond software development, all this is made possible through Skylake, the sixth-generation architecture of Core Intel processors.

SMART AND CONNECTED WORLD ...

... Or how to customize the world for the user. For example, to be able to watch in the mir-

ror and see how a garment fits you in different colors without even needing to change, as demonstrated the early virtual shopping project developed by Memomi. Another example is a project developed by Nabi, the “Baby Seat Clip” which ensures whether a child is properly secured in his seat with embedded sensors in the belt. To facilitate application developments like these, Intel has made available to developers a complete IoT platform coupled with Big Data and composed of all the essential functions (as reusable building blocks) to rapidly develop IoT-based applications: “Visualize monetize data and insight “(provides information, triggers actions, automates operations and allows the creation of services) as well as more specific functions such that “turn data into insight,” “Cloud Management “” Gateways “,” connect and things devices “, and a function “intelligence at the edge” (which filters the data, carries out data analyses and ensures that the data informs devices).

A FRAMEWORK FOR SECURING THE INTERNET OF THINGS

Intel also created a last and essential End to End Security building block to fill a crucial gap in the IoT. The market envisioned by Intel is 42 billion US dollars. Also on the security side, Intel has developed a component named Enhanced Privacy Identification, which already has two partners in its ecosystem, Microchip and Atmel. Intel provides a reference design to secure IoT applications. The results of Intel’s IoT developers program Intel launched last year are conclusive according to its CEO: 252,000 developers already use Intel’s “IOT ZONE”, and no less than 18 roadshows and 31 hackathons were held and approximately 250 co-development projects are underway.

EXTENSION OF YOU

A consequence of the fact that everything is moving towards the “Smart and Connected IT” is ultimately to extend our own capabilities. Which in turn allowed the wearables market to explode in 2015 by +170%. Which reinforces



A consequence of the fact that everything is moving towards the “Smart and Connected IT” is ultimately to extend our own capabilities. Which in turn allowed the wearables market to explode in 2015 by +170%

Intel in its partnership with a specialist in this market, watch manufacturer Fossil. Within the next few weeks co-developed connected wristwatches will be announced. Last January, at CES, was unveiled Curie, a tiny computer the size of a button. Curie among other features sports Bluetooth technology, a radio and sensors. Curie allows real-time capture of data. Intel has developed such a probe that can provide real-time data from a BMX bike like the spin, the number of flips per second, the time in seconds spent in the air, the landing impact, etc. This real-time data capture can help correct and improve very rapidly the athlete’s skills.

IOT SOFTWARE KITS...

...For a strong integration between the software and equipment on Curie platforms. To facilitate this type of development, including in digital sport, Intel has now already made available specific development kits: body IQ (physical activity) IQ and social (social interaction). Two new kits have been announced: Time IQ (managing tasks and planning) and Identity IQ (which allows to confirm an identity and provide functions when the access control). Intel has also demonstrated a proof of concept of wearables providing authentication and password: the goal is that when one approaches a device that belongs to you, it changes from locked to unlocked automatically, thanks to the “Safe” bracelet from Intel’s labs. It helps wear one’s authentication device for a PC or smartphone. The opposite is also true: when one moves away, the device automatically goes in the locked state. If one removes his bracelet, the gesture is detected

and in this case the owner must go through a traditional one or two-step authentication to reactivate the bracelet. An important detail: the technologies developed in this framework have been approved by the Bluetooth SIG and are destined to become a standard, the UAS (Bluetooth low energetic and User Authentication Service Service). For the developer, these software interfaces can be used via the Identity IQ kit that is part of the “wearable” component of the software platform available with Curie.

3D XPOINT: TOMORROW’S MEMORY AND STORAGE TO SUPPORT IOT

According to the latest projections the Internet of Things promises 50 billion connected objects by 2020, which will represent a heavy load on the current infrastructure of data centers, which could come to a crawl because of current architecture bottlenecks. Intel has partnered with Micron to develop 3D Cross Point (3D Xpoint) technology. Expected in 2016, this new class of memory technology is based on a 3D architecture designed to offer a high density and which has two possible modes of operation: as conventional memory or used in storage mode. 1000 times faster and 1000 times more enduring than NAND technology and 10 times more dense than DRAM, this new architecture promises genuine revolution in memory and storage. This technology goes on sale next year under the Intel brand Optane. Intel also plans to integrate the architecture 3D Xpoint in future Xeon processors.



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

IBM WATSON THE FUTURE PROBLEM SOLVING SUPERCOMPUTER

Watson is a question answering computing system that IBM built to apply advanced natural language processing, information retrieval, knowledge representation, automated reasoning, and machine learning technologies to the field of open domain question answering.

The key difference between QA technology and document search is that document search takes a keyword query and returns a list of documents, ranked in order of relevance to the query (often based on popularity and page ranking), while QA technology takes a question expressed in natural language, seeks to understand it in much greater detail, and returns a precise answer to the question. According to IBM, “more than 100 different techniques are used to analyze natural language, identify sources, find and generate hypotheses, find and score evidence, and merge and rank hypotheses.”

The IBM Watson system gained fame by beating human contestants on the television quiz show Jeopardy! It was also a very public trial of an advanced form of computing: a cognitive system. A system that is not simply programmed but is trained to learn, based on interactions and outcomes. A system that rivals a human’s ability to answer questions posed in natural language with speed, accura-



cy and confidence. A system that brings man and machine together. By accurately extracting facts and quickly understanding relationships in large volumes of data, the technology can help accelerate and improve decision-making across industries and in society. Watson learns from additional information, questions and mistakes to improve accuracy over time.

While today’s programmable systems can apply volumes of data to a problem, they struggle when the data is unstructured, its accuracy is uncertain, or its velocity continually accelerates. The consequence, in the end, is that decisionmaking suffers. Watson’s cognitive capabilities were designed to take on the real- world challenges of Big Data in society and across a range of industries; from the out-

TODAY IT'S ESTIMATED 80% OF THE WORLD'S DATA IS UNSTRUCTURED. INCREASINGLY, COGNITIVE SYSTEMS LIKE WATSON WILL HELP US MAKE SENSE OF IT ALL.

set, the aim was to put Watson to work, at first in healthcare and finance. Both industries confront “deluges” of unstructured data every day, and both industries have a compelling need to act on information quickly.

HEALTHCARE

In 2011, To help improve the quality of care delivered, IBM announced a pilot program with WellPoint, whose affiliated health plans cover one in nine Americans. And in March 2012, IBM launched a partnership with Memorial Sloan-Kettering Cancer Center, where work is under way to teach Watson about oncology diagnosis and treatment options. In healthcare, Watson’s natural language, hypothesis generation, and evidence-based learning capabilities allow it to function as a clinical decision support system for use by medical professionals. To aid physicians in the treatment of their patients, once a physician has posed a query to the system describing symptoms and other related factors, Watson first parses the input to identify the most important pieces of information; then mines patient data to find facts relevant to the patient’s medical and hereditary history; then examines available data sources to form and test hypotheses; and finally provides a list of individualized, confidence-scored recommendations. The sources of data that Watson uses for analysis can include treatment guidelines, electronic medical record data, notes from physicians and nurses, research materials, clinical studies, journal articles, and patient information. Des-

pite being developed and marketed as a “diagnosis and treatment advisor,” Watson has never been actually involved in the medical diagnosis process, only in assisting with identifying treatment options for patients who have already been diagnosed.

Combining its abilities to navigate the complexities of human language and to analyze massive amounts of data exceptionally quickly (over 200 million pages in 3 seconds on Jeopardy!), Watson has the potential to take advantage of new research studies, published reports and articles, as well as patient outcomes and interactions, to help physicians make evidence-based decisions. In banking Watson can pore over financial, regulatory, economic and social data across exchanges, currencies and funds at terrific speeds. For example, Citi has announced it will examine the use of Watson’s deep content analysis and evidence-based learning capabilities to advance customer interactions and to improve and simplify the experience of customers.

VISION FOR WATSON

IBM Watson is the very embodiment of this new era of cognitive systems. It represents a new class of industry specific analytic solutions that leverages deep content analysis and evidence based reasoning to accelerate and improve decisions, reduce operational costs, and optimize outcomes. It accomplishes this based on a set of transformational technologies which leverage natural language, hypothesis generation, and evidence based learning. It

combines these technologies and applies massive parallel probabilistic processing techniques to fundamentally change the way businesses look at quickly solving problems. But Watson is more than the sum of its parts; Watson is about bringing a unique combination of capabilities together in a way that's never been done before resulting in a fundamental change in the way decisions are made. IBM's vision is to define, create, and lead markets for this new class of cognitive system by:

- 1.** Addressing meaningful industry and societal challenges where conventional approaches don't work;
- 2.** Developing a cognitive class of solutions built on a secure, scalable modular framework
- 3.** Delivering demonstrable, quantifiable value as defined by the client.

A NEW ERA OF COMPUTING

This new era of computing is already changing our expectations for how technology can help us live and work in better ways. Cognitive systems can better address today's realities of explosive data growth, rapidly changing business conditions, and the need for more intuitive

interaction between people and technology. IBM's approach to the marketplace is to:

- Establish client and business value with each Watson commercial application developed
- Extend early findings in healthcare and financial services into new industries and horizontal solutions
- Embed experiences learned into go-to-market and delivery practices, partners, and IBM's Smarter Planet initiative

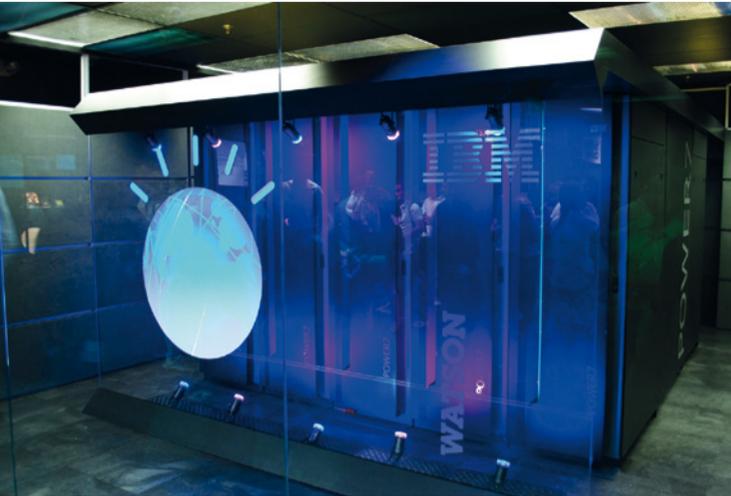
The focus is on understanding the needs of a marketplace and how best to enable transformation. This in turn will help define the methodological construct taken to delivering Watson Advisory solutions by industry and orientation as illustrated below.

WHAT'S NEXT: CURRENT AND FUTURE APPLICATIONS

Watson extends the possible ways in which organizations, institutions, and individuals will learn, process, investigate, optimize, decide and deliver. It brings into focus information and insights that were once unattainable, empowering the knowledge worker to deliver extraordinary outcomes. Watson solutions are

	HEALTHCARE	FINANCIAL SERVICES	CONTACT CENTERS
GOAL	Improve patient outcomes while dramatically reducing cost of care through personalized, evidence based medicine.	Improve knowledge worker's understanding, insights, and decision making in both institutional and retail finance as well as Insurance	Transforming the contact center from a cost-center to a strategic asset for client engagement, customer response and satisfaction through better use of information and interaction
CONSTRUCT	- Teach: Enable new methods for delivering medical training and instruction	- Research: Ask questions about an asset across a broad corpus of market information	- Market: Transform the business through with a game changing user experience across all channels
	- Practice: Enable research & delivery of personalized, evidence-based medicine	- Evaluate: Discover correlation between answers and intrinsic value of an asset.	- Sell: Automate the alignment of solutions to clients with superior client targeting and responsiveness
	- Pay: Enable the rapid evaluation and preauthorization of medical payment	- Act: Suggest what action to take in the market (i.e. market to a client, portfolio allocation)	- Support: Answers with confidence / evidence drive agent efficiency. Deflect calls to self-service for faster response

WATSON'S KEY FIGURES



Software - Watson uses IBM's DeepQA software and the Apache UIMA (Unstructured Information Management Architecture) framework. The system was written in various languages, including Java, C++, and Prolog, and runs on the SUSE Linux Enterprise Server 11 operating system using Apache Hadoop framework to provide distributed computing.

Hardware - The system is workload-optimized, integrating massively parallel POWER7 processors and built on IBM's DeepQA technology, which it uses to generate hypotheses, gather massive evidence, and analyze data. Watson employs a cluster of 90 IBM Power 750 servers, each of which uses a 3.5 GHz POWER7 eight core processor, with four threads per core. In total, the system has 2,880 POWER7 processor threads and 16 terabytes of RAM.

Speed - Watson can process 500 gigabytes, the equivalent of a million books, per second. IBM's master inventor and senior consultant Tony Pearson estimated Watson's hardware cost at about three million dollars. Its

Linpack performance stands at 80 TeraFLOPs, which is about half as fast as the cut-off line for the Top 500 Supercomputers list.

Data - The sources of information for Watson include encyclopedias, dictionaries, thesauri, newswire articles, and literary works. Watson also used databases, taxonomies, and ontologies. Specifically, DBPedia, WordNet, and Yago were used. The IBM team provided Watson with millions of documents, including dictionaries, encyclopedias, and other reference material that it could use to build its knowledge. Watson contains 200 million pages of structured and unstructured content consuming four terabytes of disk storage.

best suited to data-intensive industries and issues that:

- Require the analysis of a high volumes of both structured and unstructured data
- Benefit from the speed and accuracy of a response to a question or input provided
- Desire to systematically learn with every outcome or action taken, getting smarter with interaction and new evidence
- Have critical questions that require confidence weighted recommendations and supporting evidence

According to IBM, "The goal is to have computers start to interact in natural human terms across a range of applications and processes, understanding the questions that humans ask and providing answers that humans can understand and justify." It has been suggested by

Robert C. Weber, IBM's general counsel, that Watson may be used for legal research. The company also intends to use Watson in other information-intensive fields, such as telecommunications, financial services, and government. Watson is based on commercially available IBM Power 750 servers that have been marketed since February 2010. IBM also intends to market the DeepQA software to large corporations, with a price in the millions of dollars, reflecting the \$1 million needed to acquire a server that meets the minimum system requirement to operate Watson. IBM expects the price to decrease substantially within a decade as the technology improves.

In 2013, it was reported that three companies were working with IBM to create apps embedded with Watson technology. Fluid is developing an app for retailers, one called

ACCORDING TO IBM, “MORE THAN 100 DIFFERENT TECHNIQUES ARE USED TO ANALYZE NATURAL LANGUAGE, IDENTIFY SOURCES, FIND AND GENERATE HYPOTHESES, FIND AND SCORE EVIDENCE, AND MERGE AND RANK HYPOTHESES.”

“The North Face”, which is designed to provide advice to online shoppers. Welltok is developing an app designed to give people advice on ways to engage in activities to improve their health. MD Buyline is developing an app for the purpose of advising medical institutions on equipment procurement decisions.

In November, 2013, IBM announced it would make Watson’s API available to software application providers, enabling them to build apps and services that are embedded with Watson’s capabilities. To build out its base of partners who create applications on the Watson platform, IBM consults with a network of venture capital firms, which advise IBM on which of their portfolio companies may be a logical fit for what IBM calls the Watson Ecosystem. Thus far, roughly 800 organizations and individuals have signed up with IBM, with interest in creating applications that could use the Watson platform.

On February 6, 2014, IBM announced plans to invest \$100 million in a 10-year initiative to use Watson and other IBM technologies to help countries in Africa address development problems, beginning with healthcare and education.

On June 3, 2014, three new Watson Ecosystem partners were chosen from more than 400 business concepts submitted by teams spanning 18 industries from 43 countries. “These bright and enterprising organizations have discovered innovative ways to apply Watson that can deliver demonstrable business bene-

fits,” said Steve Gold, vice president, IBM Watson Group. The winners are Majestyk Apps with their adaptive educational platform, FANG (Friendly Anthropomorphic Networked Genome); Red Ant with their retail sales trainer; and GenieMD with their medical recommendation service.

On July 9, 2014, Genesys Telecommunications Laboratories announced plans to integrate Watson to improve their customer experience platform, citing the sheer volume of customer data to analyze is staggering. Watson has also been integrated with databases including Bon Appétit magazine to perform a recipe generating platform. Besides, Watson is being used by Decibel, a music discovery startup, in its app MusicGeek which uses the supercomputer to provide music recommendations to its users. The use of the artificial intelligence of Watson has also been found in hospitality industry. GoMoment uses Watson for its Rev1 app, which gives hotel staff a way to quickly respond to questions from guests. Arria NLG has built an app that helps energy companies stay within regulatory guidelines, making it easier for managers to make sense of thousands of pages of legal and technical jargon.

Where, how, and when to put Watson to work may be one of the greatest challenges in the years to come. We have only begun to scratch the surface of what is possible in this new era of computing.



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

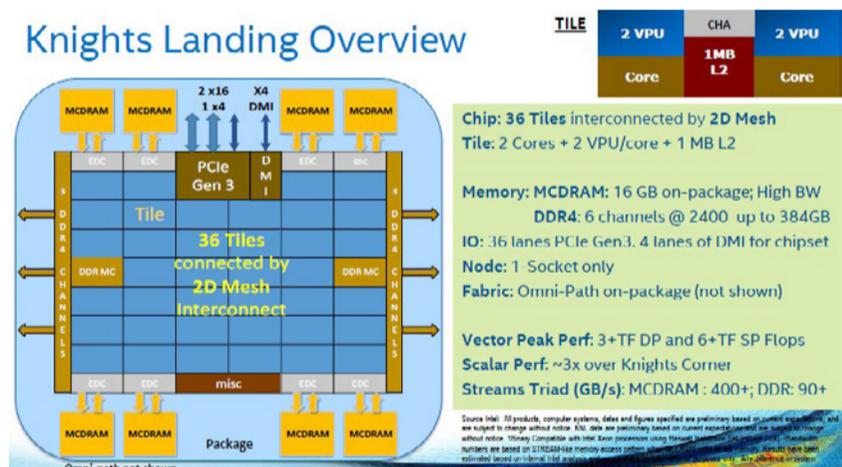
INSIDE INTEL KNIGHTS LANDING ARCHITECTURE

Knights Landing is the codename for Intel's forthcoming Xeon Phi processor dedicated to HPC. Few details have emerged until now, besides that it will sport 72 cores, be able to do 240 threads in parallel and have a massive 8-billion transistor die. However, several key announcements have emerged since SC'15 in november, and the picture is getting more clear while the launch date of Q1 2016 approaches.

The 14nm successor to Knights Corner (1st gen Xeon Phi), Knights Landing implements AVX-512, Multi-Channel DRAM (MCDRAM), and a new CPU core based on Intel's Silvermont architecture. Knights Landing is now shipping to Intel's first customers and developers as part of their early ship program, and pre-production systems for demonstrating supercomputer designs are up and running. Knights Landing is ultimately ramping up for general availability in Q1 of 2016.

ONE PLATFORM, THREE PROCESSORS

First things first : instead of one product, Knights Landing will in fact, be launched as three different products, packaged differently to suit different usage scenarios. Given that the Knights Landing implementation is available as a standalone processor as well as a coprocessor, you



might be thinking that Intel expects for a lot of the machinery built using Knights Landing will be a mix of Xeon and Xeon Phi systems clustered together and working side-by-side but not with the Xeon Phi being linked as a coprocessor to the Xeons

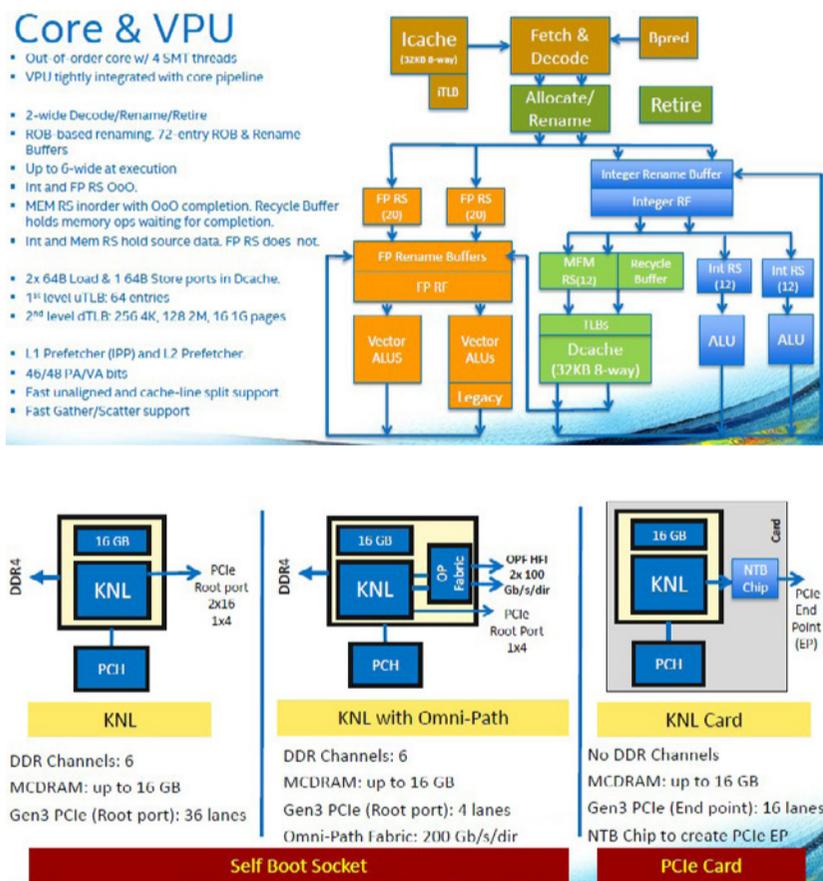
The base bootable Knights Landing chip has 16 GB of MCDRAM high bandwidth memory right on the package and DDR4 memory controllers to link out to a maximum of 384 GB of regular DRAM memory. The chip has two PCI-Express 3.0 x16 ports and one x4 root port as well as a south-bridge chipset to link to various I/O devices. The

KNIGHTS LANDING EARLY PERFORMANCE FIGURES : VERY PROMISING

Across all 72 cores, the Knights Landing chip will deliver more than 6 teraflops of single precision and more than 3 teraflops of single precision performance thanks to the two AVX512 vector processing units that are welded onto each core. (The can process 32 single precision and 16 double precision operations per clock per tile.) The Knights Landing chip

is the first of Intel's Xeon family that will have support for this AVX512 vector processing unit. Intel showed some performance benchmarks that compared a single Xeon Phi with 72 cores testing against a two-socket Xeon E5 v3 server using ten-core E5-2697 processors on a variety of raw processor and application workload benchmarks. The single

Xeon Phi has about 2.5X the peak raw double-precision teraflops as a pair of Xeon E5s, and also can run the AlexNet neural network training algorithm about 2.5X as fast and the STREAM memory bandwidth test about 3.5X as fast. If you adjust this for performance per watt, the gap is even larger.



chip on the card, linking to the processor through one PCI-Express x16 port, as a PCI-Express endpoint. This version of the chip will also have its DDR4 memory channels deactivated and will only have the 16 GB of MCDRAM near memory for the processor to access. It is unclear if all of the memory modes developed for Knights Landing will be supported in the coprocessor, but given the latencies of moving data over the PCI-Express bus, we suspect not. If the price on the free-standing Knights Landing cards is small enough, it is possible that these sell well, but we expect most enterprises, hyperscalers, and HPC centers will be interested in the self-booting versions of the Knights Landing chips, not the coprocessors.

Omni-Path interconnect is not on the chip itself, but is rather implemented in the package, with each PCI-Express x16 port having its own bi-directional, 100 Gb/sec ports. Each port can deliver 25 GB/sec of bandwidth in both directions. Integrating an InfiniBand or Omni-Path interconnect directly on a Xeon or Xeon Phi die is tricky, so Intel is integrating on the package first as a means of lowering the overall cost and power consumption in the Knights Landing platform.

The free-standing Knights Landing co-processor card will put a non-transparent bridge (NTB)

INTEL OMNI-PATH ARCHITECTURE

Meanwhile Knights Landing's partner in processing, Intel's Omni-Path Architecture, has formally been launched at SC15. Intel's own take on a high bandwidth low-latency interconnect for HPC, Omni-Path marks Intel's greatest efforts yet to diverge from InfiniBand and go their own way in the market for interconnect fabrics.

INTEL SYSTEM SCALABLE FRAMEWORK

Ultimately Knights Landing and Omni-Path Architecture are part of Intel's larger efforts to build a whole ecosystem, which they've been calling the System Scalable Framework.



TECH ZONE

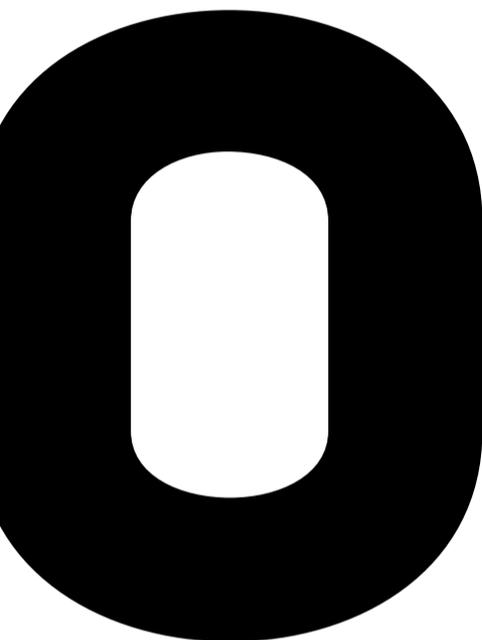


HPC
REVIEW

CERN SEASON 2

NEW FRONTIERS
IN PHYSICS AND
EXASCALE IN 2020

In early 2013, the Large Hadron Collider (LHC) shut down for two years of planned maintenance and repairs. Now the world's most powerful accelerator is ready to start up again, this time at 13 TeV – almost double its previous energy. This new energy frontier will allow researchers to probe new boundaries in our understanding of the fundamental structure of matter.



On 4 July 2012, the ATLAS and CMS experiments at CERN announced the discovery of a Higgs boson, a particle with a mass of 126 GeV. The Higgs is the simplest manifestation of the Brout-Englert-Higgs mechanism that gives particles mass. It is the final particle in the Standard Model – the theory that explains the

fundamental particles and the forces between them – to be experimentally verified. Increasing the energy of the LHC will increase the chance of creating Higgs bosons in collisions, which means more opportunity for researchers to measure the Higgs precisely and to probe its rarer decays. High-energy collisions could detect small, subtle differences between what the boson looks like in experiments, and what the Standard Model predicts.

EXOTIC PARTICLES

Some theories predict that there could be a whole new set of particles out there that physicists cannot detect because they don't interact through the electromagnetic force. But if these “dark sector” particles have mass, they will interact with the field associated with the Higgs boson. So the Higgs boson becomes a contact point between the Standard Model and new, more exotic particles. Dark matter Invisible dark matter makes up most of the universe – but we can only detect it from its

gravitational effects. But what is dark matter? One idea is that it could contain “supersymmetric particles” – hypothesized particles that are partners to those already known in the Standard Model. The data from higher-energy running at the LHC could provide more direct clues to resolve this mystery.

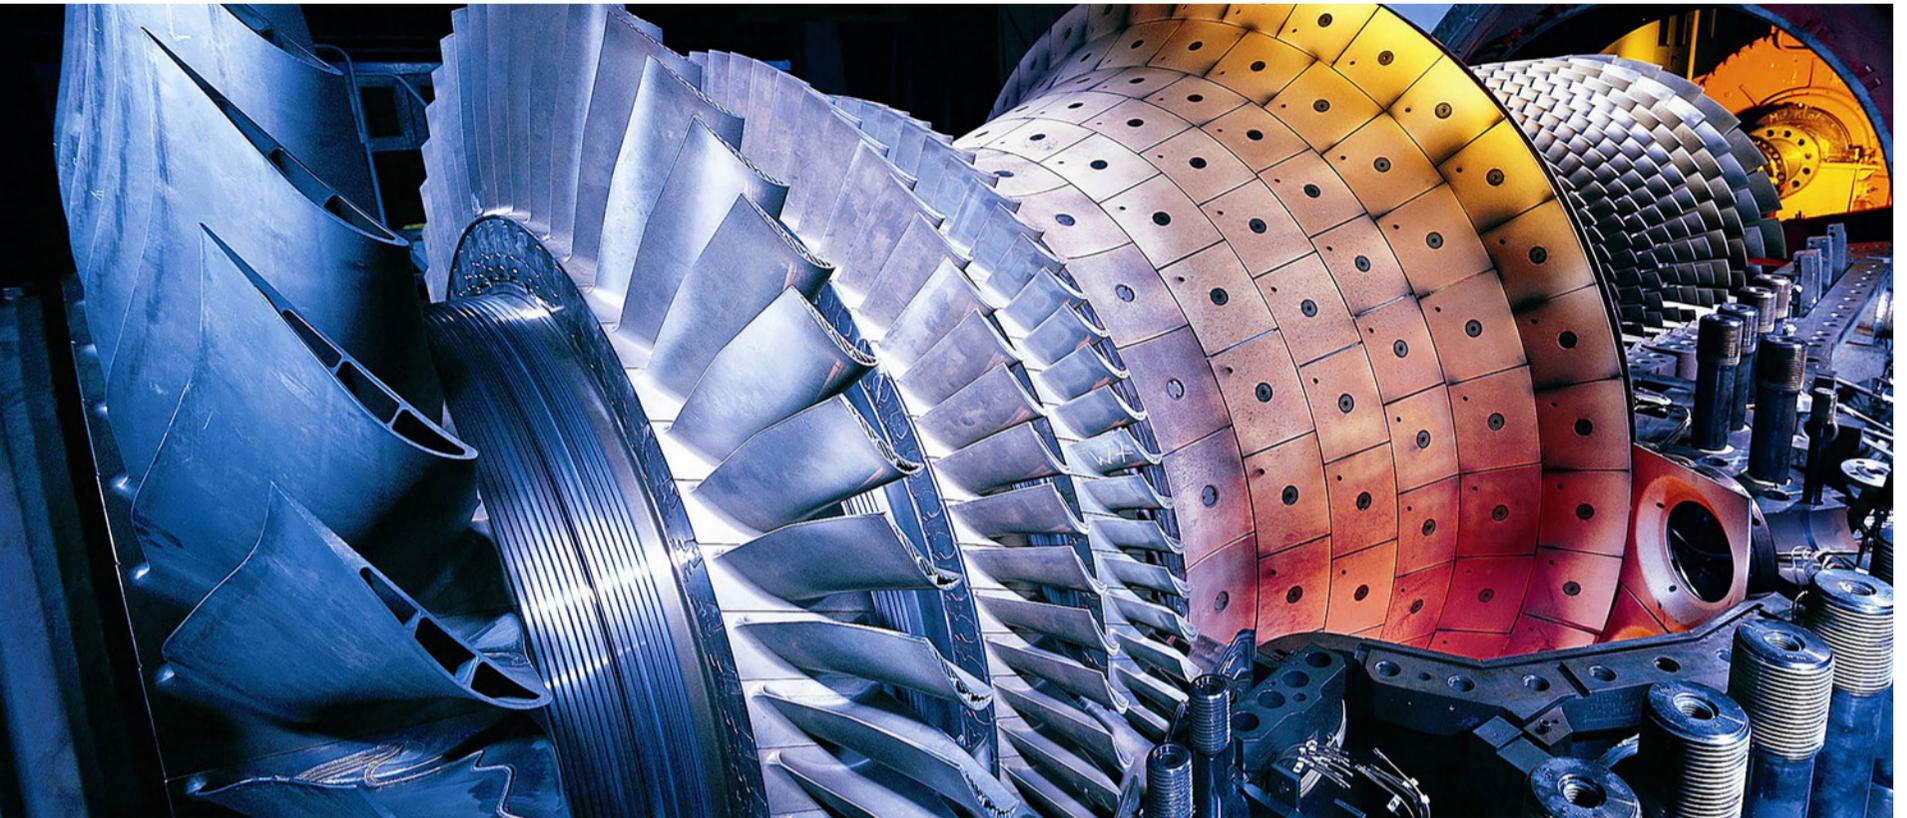
SUPERSYMMETRY

The Standard Model has worked beautifully to predict what experiments have shown so far about the basic building blocks of matter, but the theory is incomplete. Supersymmetry is an extension of the Standard Model that aims to fill some of the gaps. It predicts a partner particle for each particle in the Standard Model. These new particles would solve a major problem, fixing the mass of the Higgs boson. If the theory is correct, supersymmetric particles should appear in high-energy collisions at the LHC.

EXTRA DIMENSIONS

Why is gravity so much weaker than the other fundamental forces? Perhaps we don't feel the full effect of gravity because part of it spreads to extra dimensions. But how could we test for extra dimensions? One option is to find evidence of particles that can exist only if extra dimensions are real. Theories that require extra dimensions predict that, just as atoms have a low-energy ground state and excited high-energy states, there would be heavier versions of standard particles in other dimensions. Such heavy particles could be revealed at the high energies the LHC will reach in Run 2.

CERN RUN 2: FACTS AND FIGURES



Circumference 26 659 m

Dipole operating temperature
1.9 K (-271.3°C)

Number of magnets 9593

Number of main dipoles 1232

Number of main quadrupoles 392

Number of RF cavities

8 per beam

Nominal energy, protons
6.5 TeV

Nominal energy, ions
2.56 TeV/u (energy
per nucleon)

Nominal energy, protons collisions
13 TeV

No. of bunches per proton beam
2808

No. of protons per bunch (at start)
 1.2×10^{11}

Number of turns per second
11245

Number of collisions per second
1 billion

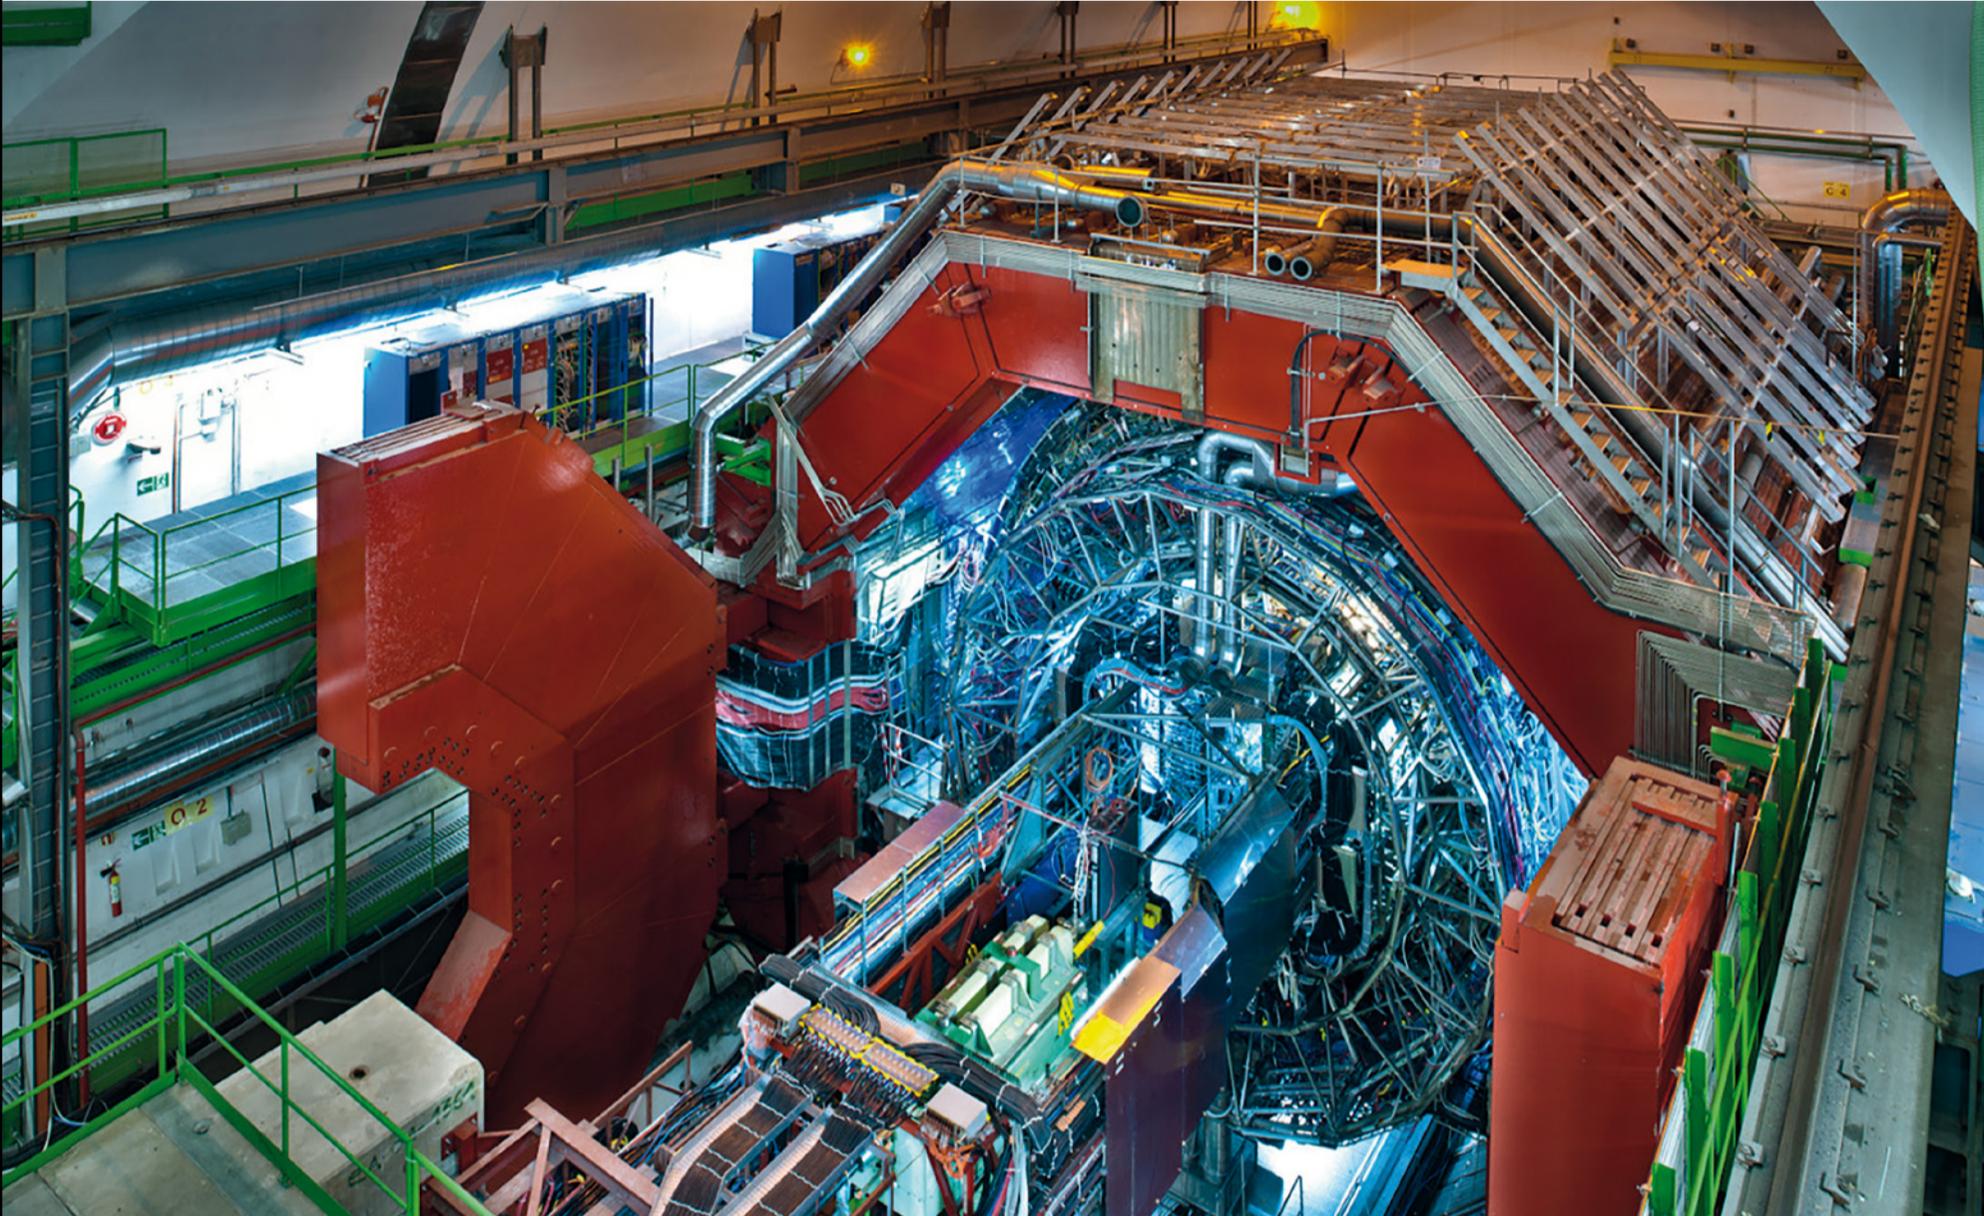
ANTIMATTER

Every particle of matter has a corresponding antiparticle, exactly matching the particle but with opposite charge. For the electron, for example, there is an “antielectron” called the positron – identical in every way but with a positive electric charge. But when matter and antimatter come into contact, they annihilate, disappearing in a flash of energy. The Big Bang should have created equal amounts of matter and antimatter. So why is there far more matter than antimatter in the universe? Running at higher energy will allow the production of more antiparticles for CERN’s antimatter programme – helping physicists to check if the properties of antimatter differ from those of matter.

QUARK–GLUON PLASMA

For a few millionths of a second, shortly after the Big Bang, the universe was filled with an astonishingly hot, dense soup made of all kinds of particles moving at near light speed. This mixture was dominated by quarks – fundamental bits of matter – and by gluons, carriers of the strong force that normally “glue” quarks together into familiar protons, neutrons and other species. In those first evanescent moments of extreme temperature, however, quarks and gluons were bound only weakly, free to move on their own in what’s called a quark–gluon plasma. The higher energy collisions at the LHC will allow new and more detailed characterization of this quark–gluon plasma.

LES GRANDES ÉVOLUTIONS D'INFRASTRUCTURE



ALICE

THE DETECTOR

The heavy-ion detector ALICE studies the quark-gluon plasma, a phase of matter that formed just after the Big Bang and that gave rise to the particles that make up the matter in the universe today. A droplet of this primordial matter is produced when very high-energy lead nuclei collide at the LHC.

LONG SHUTDOWN 1

Many hundreds of members of the ALICE collaboration worked together to improve the detector. Among the consolidation and improvements to its 19 subdetectors, they installed a new calorimeter (DCAL) extending the range covered by the electromagnetic calorimeter (EMCAL). The EMCAL now covers a wider angle, allowing the measurement of the energy of the photons and electrons over a larger area. The TRD (Transition Radiation Detector) that detects

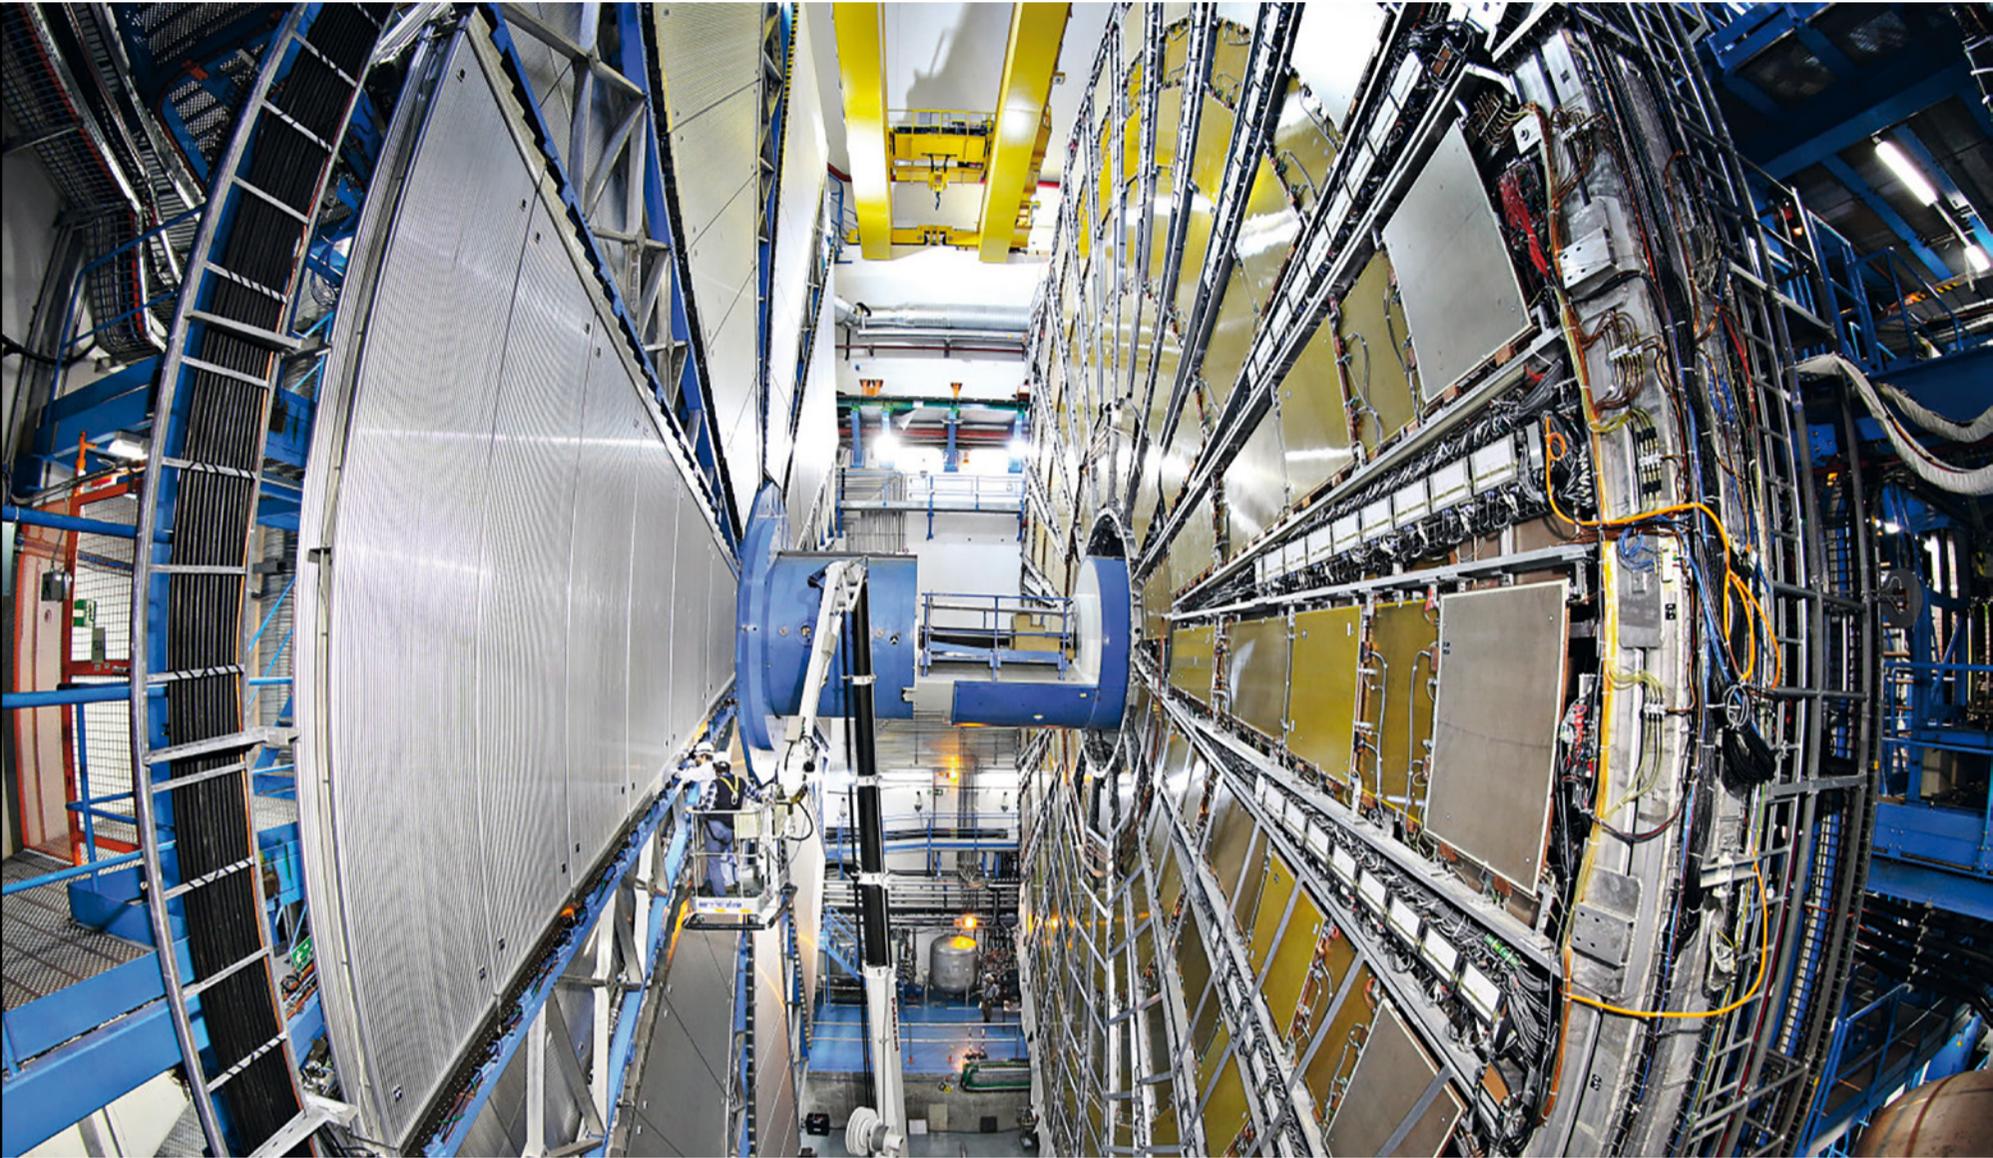
particle tracks and identifies electrons has also been completed with the addition of 5 more modules.

PHYSICS AT 13 TEV

During Run 2, the ALICE collaboration will further explore the properties of quark-gluon plasma. The higher collision energy will lead to more high-energy particles interacting with this hot medium. With more collisions, ALICE will see an increase in statistics that will allow more detailed measurements.

THE DETECTOR: 10,000 tons, 26 meters long, 16 meters high, 16 meters wide

THE COLLABORATION: 1500 members from 154 institutes in 37 countries



ATLAS & CMS

The ATLAS and CMS detectors have a broad physics programme including investigating the recently discovered Higgs Boson, searching for extra dimensions and particles that could make up dark matter, as well as continuing systematic studies of the Standard Model.

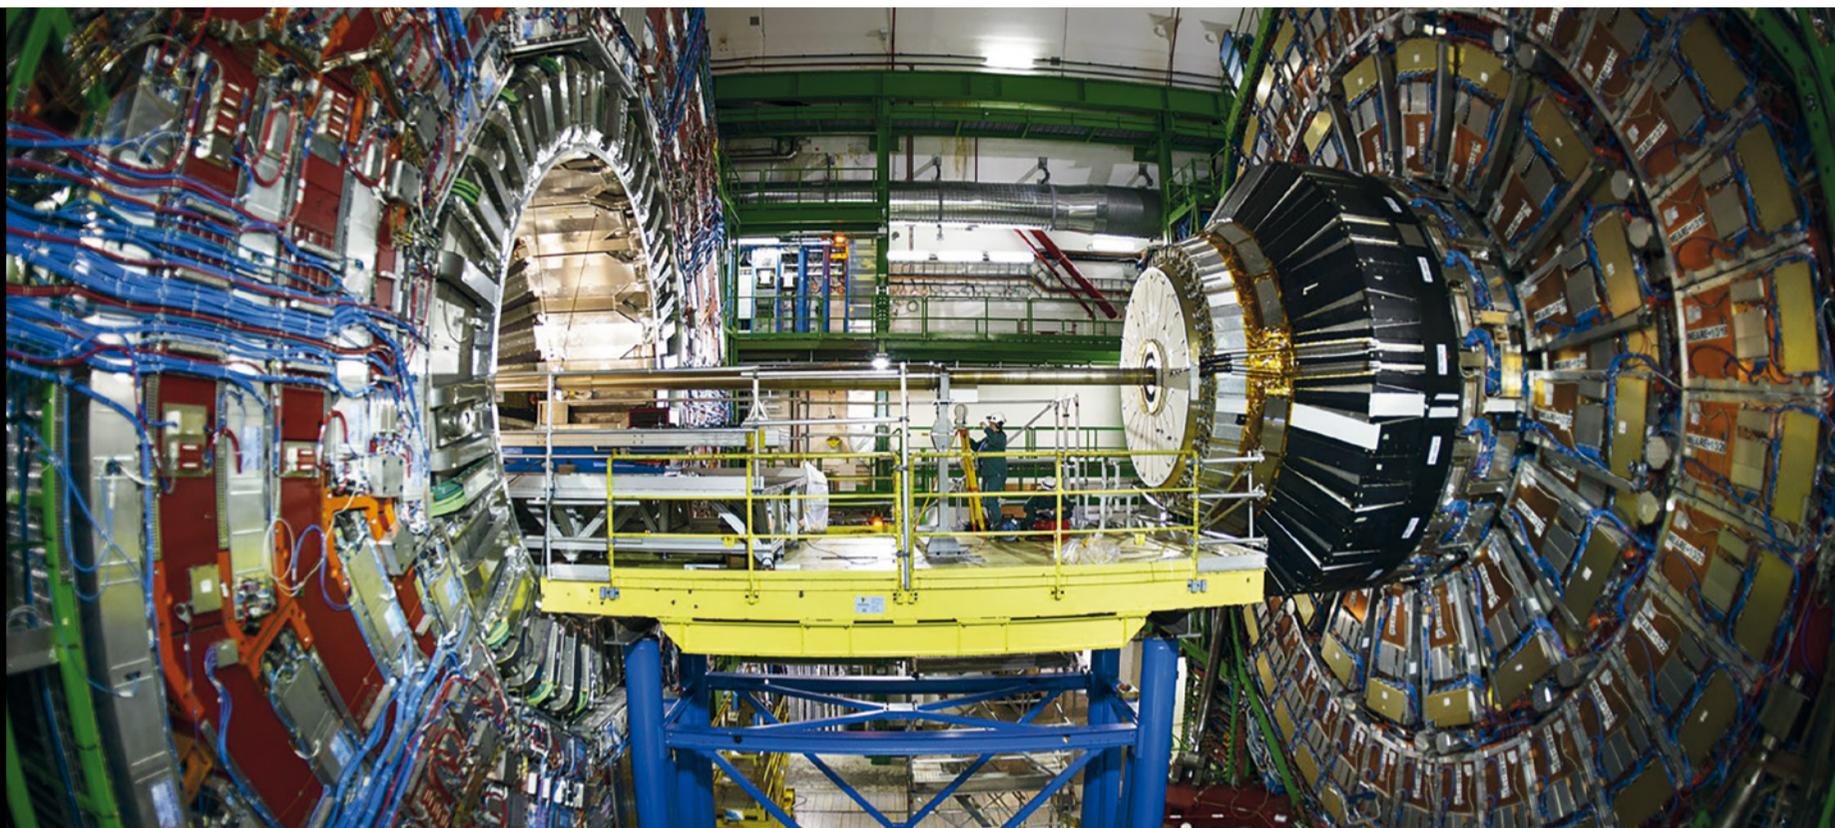
ATLAS: LONG SHUTDOWN 1

In Run 2 the number of collisions per beam crossing will increase. The ATLAS pixel detector was improved with the insertion of a fourth and innermost layer that will provide the experiment

with better vertex identification, essential to distinguish interesting collisions. The collaboration also used the shutdown to improve the general ATLAS infrastructure, including electrical power, cryogenic and cooling systems. The gas system of the TRT, which contributes to the identification of electrons as well as track reconstruction, was modified significantly to minimize losses. New chambers were added in the muon spectrometer and the calorimeter readout consolidated. The forward detectors were upgraded to provide a better measurement of the LHC

↑ **ATLAS**
THE DETECTOR: 7000 tons, 46 meters long, 25 meters high, 25 meters wide
THE COLLABORATION: 3000 members from 174 institutes in 38 countries

luminosity, and a new aluminium beam pipe was installed to reduce the background. Another challenge is the increased collision rate. The whole detector readout system was improved to be able to run at 100 KHz and all data acquisition software and monitoring applications were re-engineered. The trigger system was redesigned, going from 3 levels to 2 while implementing smarter and faster selection



algorithms. The time needed to reconstruct ATLAS events also needed to be reduced, even with more activity in the detector. A very ambitious upgrade of simulation, reconstruction and analysis software was completed, and a new generation of data management tools on the GRID was implemented.

CMS: LONG SHUTDOWN 1

The biggest priority was to mitigate the effects of radiation on the performance of the Tracker, by equipping it to operate at low temperatures (down to -20°C). This required changes to the cooling plant, and extensive work on the environment control of detector and cooling distribution to prevent condensation or icing. The central beam-pipe was replaced by a narrower one in preparation for the installation in 2016-'17 of a new Pixel Tracker that will better measure the momenta and points of origin of charged particles. A fourth measuring station was

added to each muon endcap, in order to maintain discrimination between low-momentum muons and background as the LHC beam intensities increase. Complementary to this was the installation at each end of the detector of a 125-tonne composite shielding wall to reduce neutron backgrounds. A luminosity-measuring device, the Pixel Luminosity Telescope, was installed on either side of the collision point around the beam-pipe. Other major activities included replacing photo-detectors in the hadron calorimeter with better-performing designs, moving the muon readout to more accessible locations for maintenance, installation of the first stage of a new hardware triggering system and consolidation of the solenoid's magnet cryogenic system and of the power distribution. The software and computing systems underwent a significant overhaul during the shutdown to reduce the

↑ CMS THE DETECTOR:

14000 tons, 21 meters long, 15 meters high, 15 meters wide

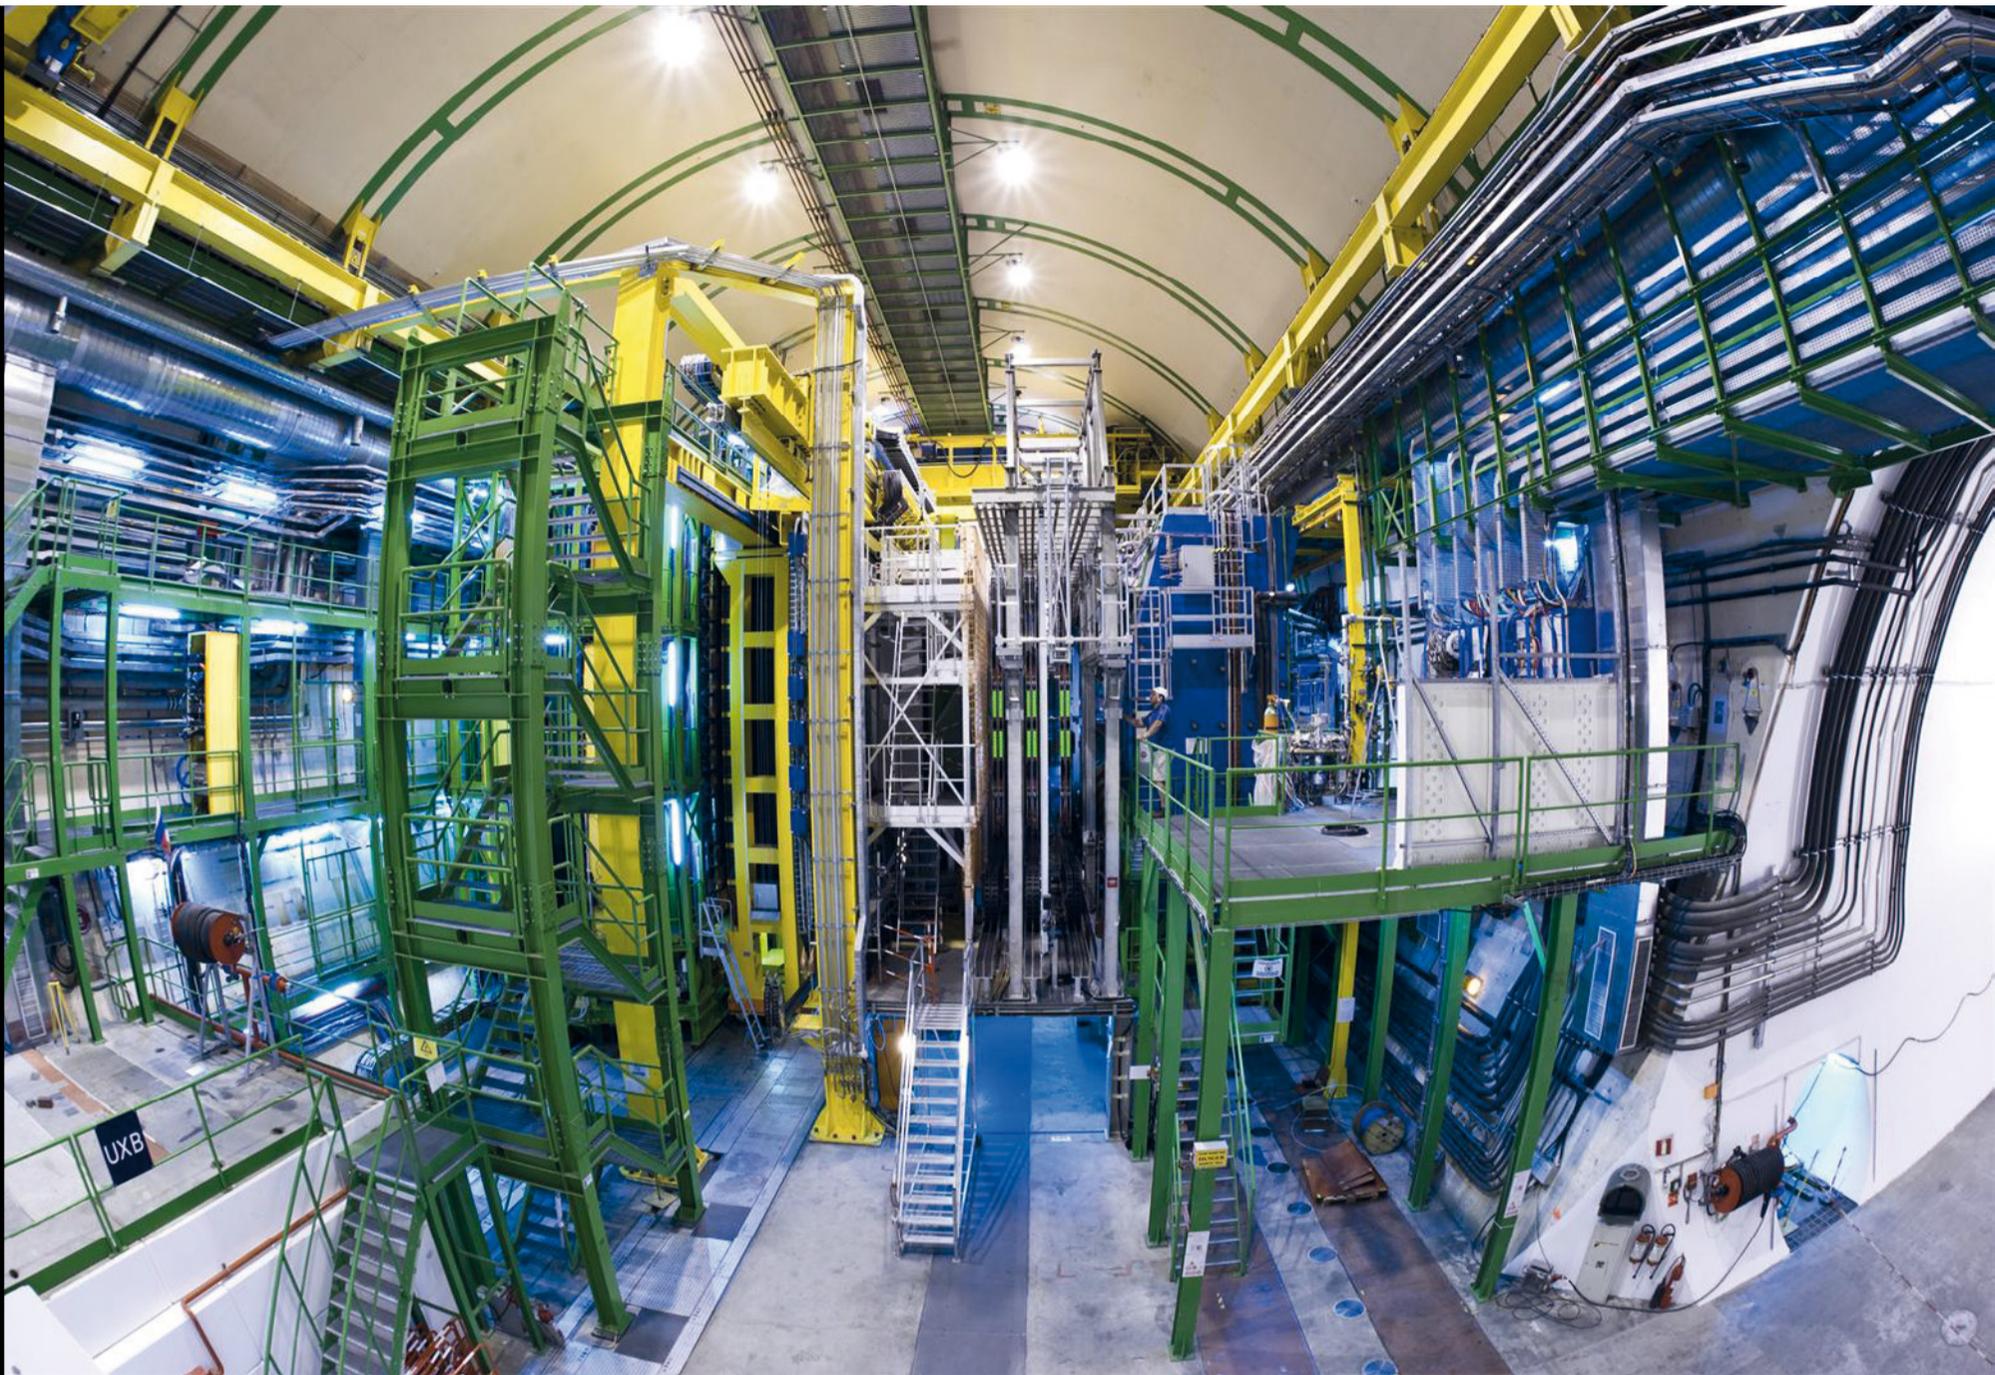
THE COLLABORATION:

3000 members from 185 institutes in 42 countries

time needed to produce analysis datasets.

ATLAS & CMS: PHYSICS AT 13TEV

The second run of the LHC will allow ATLAS and CMS physicists to extend the search for physics beyond the Standard Model. The increase in energy opens up unexplored regions for the production of high-mass new particles. Standard Model physics studies, in particular in the Higgs sector, will largely profit from the increased rate of collisions that will allow the physicists to improve the precision of our measurements.



LHCb

The LHCb experiment investigates the slight differences between properties of matter and antimatter by studying a type of particle containing the “beauty quark” (or “b quark”). Analysing rare beauty particle decays allows LHCb to investigate a possible manifestation of new physics effects.

LONG SHUTDOWN 1

To make the most of the high-energy collisions, the LHCb

collaboration improved its detector and facilities. They installed the new HeRScheL detector to distinguish rare processes in which particles are observed in the detector but not along the beampipe. One section of the beryllium beam pipe was replaced. In addition the new beam pipe support structure is now much lighter.

PHYSICS AT 13TEV

LHCb published many results

↑ LHCb THE DETECTOR:

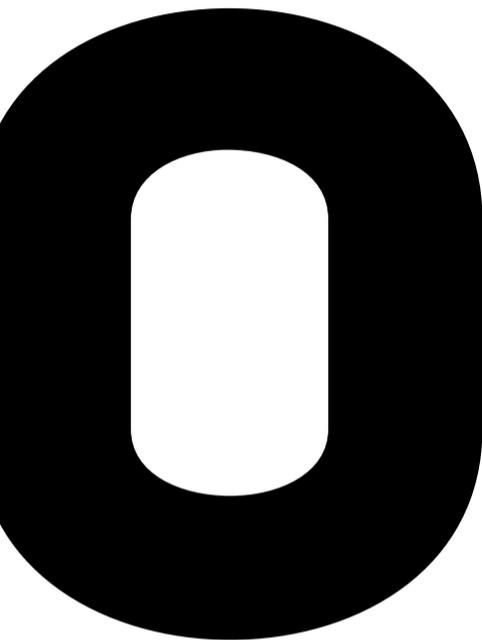
5600 tons, 21 meters long, 11 meters high, 13 meters wide

THE COLLABORATION:

1128 members from 68 institutes in 16 countries

based on data collected during the Large Hadron Collider’s first three-year run. But this was only a beginning. Collisions at 13 TeV will double the production rates of beauty hadrons enabling LHCb to get even more precise, interesting and hopefully surprising results.

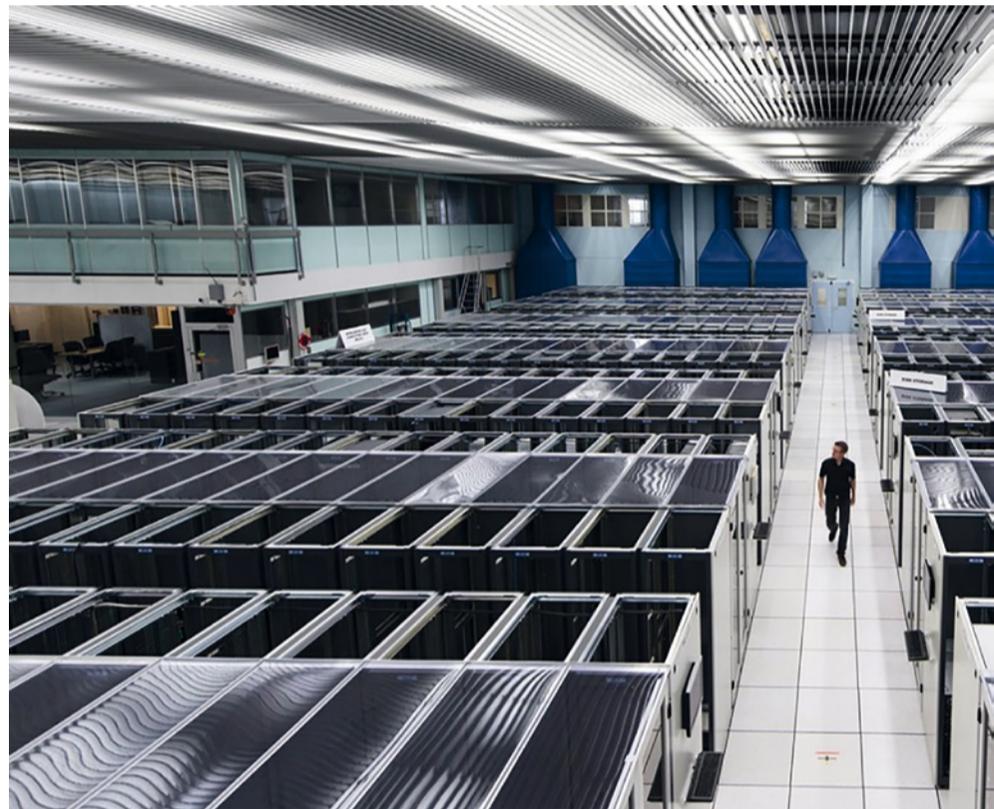
2020 : TARGETING EXASCALE



originally, the LHCb detector was designed for operation with moderate luminosity and low pile-up. However, in 2010, the collaboration opted for “luminosity levelling”, a novel solution which allows the experiment to adapt automatically to normal variations in luminosity which occur during an LHC run. In

this way, the detector operates optimally at all times. “For the second LHC run, we will have to redefine the luminosity, to adapt to the conditions at the new energy of 13 TeV,” explains Patrick Koppenburg, physics coordinator of the LHCb collaboration. “However, the most important experimental challenge for us will be the new trigger system.”

The trigger rapidly sorts the most interesting data from the data that can be discarded without a significant loss of information. The zero-level (i.e. the first level) trigger system of LHCb “only” lets through one sixteenth of the initial data, but even that is too much to be stored permanently. “During the first run, the data which got through the zero-level trigger was treated using the Hlt1 and Hlt2 algorithms, almost in real-time,” says Patrick Koppenburg. “Then we observed that on certain crucial parameters such as lifetime acceptance, differences started appearing between the values of the trigger system and those calculated after analysis. So from January 2015, we will temporarily store to disk all the data filtered by Hlt1. Then, we’ll run Hlt2 on the



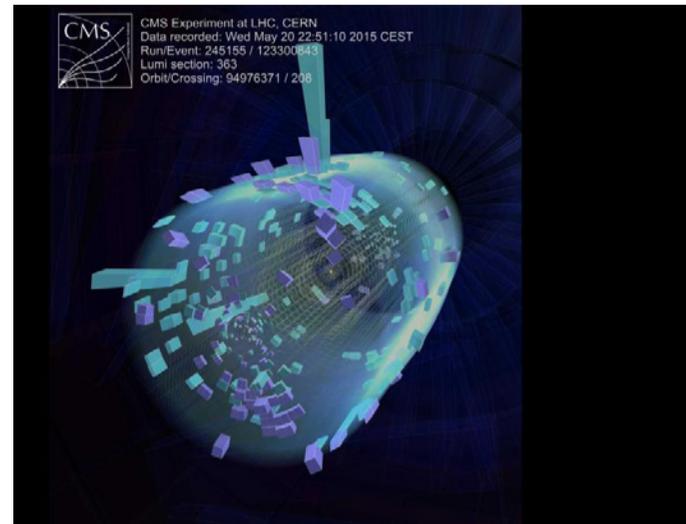
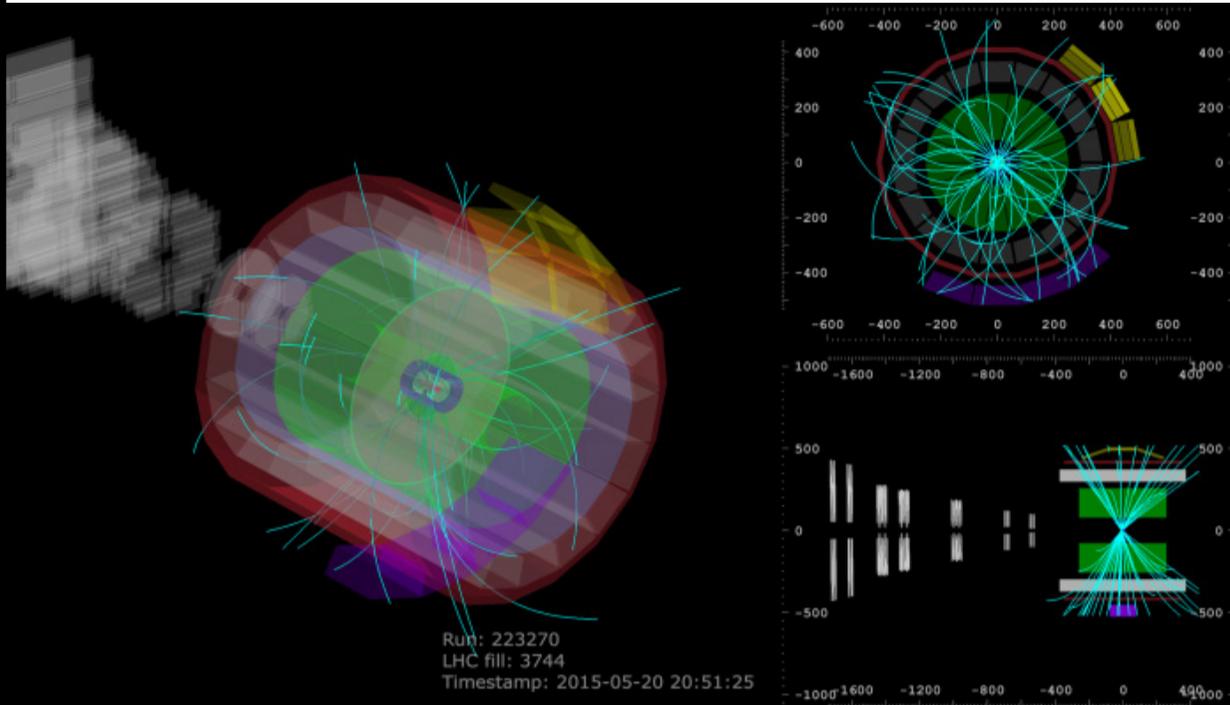
data after having calibrated the detector. This procedure will enable us to eliminate most discrepancies at source, as the reconstruction of the data will be the same in the trigger system as in the final analysis.”

The expectations of the LHCb collaboration for the second LHC run are focused on two topics: “In the first months, running at low luminosity, we’ll do cross-section measurements at 13 TeV of the production of charm, the B particle, and the Z and W particles and, more generally, we’ll measure the charged forward multiplicity,” Patrick Koppenburg concludes. “Then we will continue to accumulate statistical data for our precision studies of b and charm physics.”

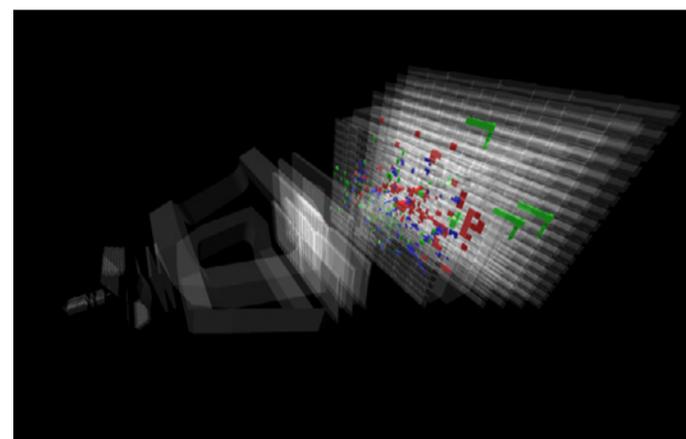
PERFORMANCE IMPROVEMENTS

The experiments at the Large Hadron Collider (LHC) will start taking data at the new energy frontier of 13 teraelectronvolts (TeV) - nearly

COLLISIONS AT 13 TEV: FIRST PICTURES



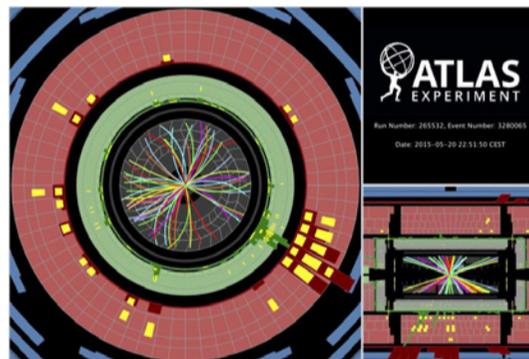
... through the CMS detector



... and through the LHCb detector

Protons collided in the Large Hadron Collider (LHC) at the record-breaking energy of 13 teraelectronvolts (TeV) for the first time. These test collisions were to set up systems that protect the machine and detectors from particles that stray from the edges of the beam. A key part of the process was the set-up of the collimators. These devices which absorb stray particles were adjusted in colliding-beam conditions. This set-up will give the accelerator team the data they need to ensure that the LHC magnets and detectors are fully protected. Today the tests

Protons collide at 13 TeV sending showers of particles through the ALICE detector...



... through the ATLAS detector

continue. Colliding beams will stay in the LHC for several hours. The LHC Operations team will continue to monitor beam quality and optimisation of the set-up.

double the energy of collisions in the LHC's first three-year run. These collisions, which will occur up to 1 billion times every second, will send showers of particles through the detectors.

With every second of run-time, gigabytes of data will come pouring into the CERN Data Centre to be stored, sorted and shared with physicists worldwide. To cope with this mas-

sive influx of Run 2 data, the CERN computing teams focused on three areas: speed, capacity and reliability.

"During Run 1, we were storing 1 gigabyte-per-second, with the occasional peak of 6 gigabytes-per-second," says Alberto Pace, who leads the Data and Storage Services group within the IT Department. "For Run 2, what was once our "peak" will now be considered ave-

THE TRANSITION TO SEASON 2 HIGHLIGHTED NEW NEEDS FOR THE TREATMENT AND CONSERVATION OF THE DATA.

rage, and we believe we could even go up to 10 gigabytes-per-second if needed.”

At CERN, most of the data is archived on magnetic tape using the CERN Advanced Storage system (CASTOR) and the rest is stored on the EOS disk pool system – a system optimized for fast analysis access by many concurrent users. Magnetic tapes may be seen as an old-fashioned technology. They are actually a robust storage material, able to store huge volumes of data and thus ideal for long-term preservation. The computing teams have improved the software of the tape storage system CASTOR, allowing CERN’s tape drives and libraries to be used more efficiently, with no lag times or delays. This allows the Data Centre to increase the rate of data that can be moved to tape and read back.

REDUCING THE RISK OF DATA LOSS

Reducing the risk of data loss - and the massive storage burden associated with this - was another challenge to address for Run 2. The computing teams introduced a data ‘chunking’ option in the EOS storage disk system. This splits the data into segments and enables recently acquired data to be kept on disk for quick access. “This allowed our online total data capacity to be increased significantly,” Pace continues. “We have 140 petabytes of raw disk space available for Run 2 data, divided between the CERN Data Centre and the Wigner Data Centre in Budapest, Hungary. This translates to about 60 petabytes of storage, including back-up files.” 140 petabytes (which is equal to 140 million gigabytes) is a very large number indeed - equivalent to over a thousand of full HD-quality movies.

DATA SLICING AND REPLICATION

In addition to the regular “replication” approach - whereby a duplicated copy is kept for

all data - experiments will now have an option to scatter the data across multiple disks. This “chunking” approach breaks the data into pieces. Use of reconstruction algorithms means that content will not be lost even if multiple disks fail. This not only decreases the probability of data loss, but also cuts in half the space needed for back-up storage. Finally, the EOS system has also been further improved to achieve the goal of more than 99.5% availability for the duration of Run 2. From quicker storage speeds to new storage solutions, CERN is well-prepared for all of the fantastic challenges of Run 2.

STORAGE IMPROVEMENTS FOR FUTURE NEEDS

The transition to season 2 highlighted new needs for the treatment and conservation of the data. This implies that 1 billion collisions occur each second, generating avalanches of particles in the detectors. Every second of operation of the LHC and its experiments, generates several gigabytes of data which need to be stored, sorted and shared with physicists worldwide. In order to cope with this massive influx of data, CERN’s Data Storage group focused on three aspects: speed, capacity and reliability. “During the first operating period, we stored 1 gigabyte per second, with occasional peaks of 6Gb / s says Alberto Pace, who runs the Storage group data and service data within CERN’s IT department. For the second operation period, 6 GB / s is now considered an average, and we believe we can go up to 10 GB / s if necessary. “At CERN, the bulk of the data is stored on magnetic tape through our CASTOR improved storage system and the rest is stored on EOS, our common storage discs with an optimized filesystem in order to allow quick access to a large number of concurrent users.

CERN'S INFRASTRUCTURE HAS ALSO EVOLVED DURING THE SEASON 2 OF THE PARTICLE ACCELERATOR.

Tape technology may seem exceeded. Yet it is a robust storage solution for recording huge amounts of data and thus ideal for long-term preservation. The teams improved the CASTOR storage system software. The reelstands and magnetic tape libraries can thus be used more effectively, without downtime or delay. The centre can thus increase the pace the data is transferred to and read from tape.

ROADMAP DEVELOPMENTS AND EXASCALE IN 2020

CERN's infrastructure has also evolved during the Season 2 of the particle accelerator. Dirk Düllmann, deputy head of group data and storage services in the IT Department at CERN, detailed in the month of June 2015 roadmap to achieve developments the goal of exascale 2020. Dirk Düllmann provides storage services and develops data management frameworks for the physics community at CERN. It is responsible for the development and evolution CERN storage components and high performance disk pools for analyzing LHC data. Projected data generated by CERN is slated to be multiplied by 14 by 2020, for a grand total of over 200 Peta-bytes. As a comparison, Google's search engine handles 100 Petabytes per year, and Facebook 180 Petabytes per year.

FROM 2020 TO 2030 : AFTER THE LHC, THE HL-LHC

The current phase, dubbed Season 2, extends until the end of 2017, with a new hiatus planned in 2018, before resuming operations 3 for season 2019 to 2021. 2022 will a pivotal year, as developments planned for the LHC will evolve towards the HL-LHC (High Luminosity LHC) phase which is planned to last until 2030. These developments go hand in hand with a

substantial increase in energy needs, since it will increase from 8 megawatts 20 megawatts. System memory will be increased to 64 Peta-bytes, and the LHC's storage capacity should reach a staggering 1 exabyte (1000 terabytes). Despite these quantitatively significant developments, the availability will be multiplied by 7, with a MTTI (mean time to interrupt) reduced to 1 day instead of 7.

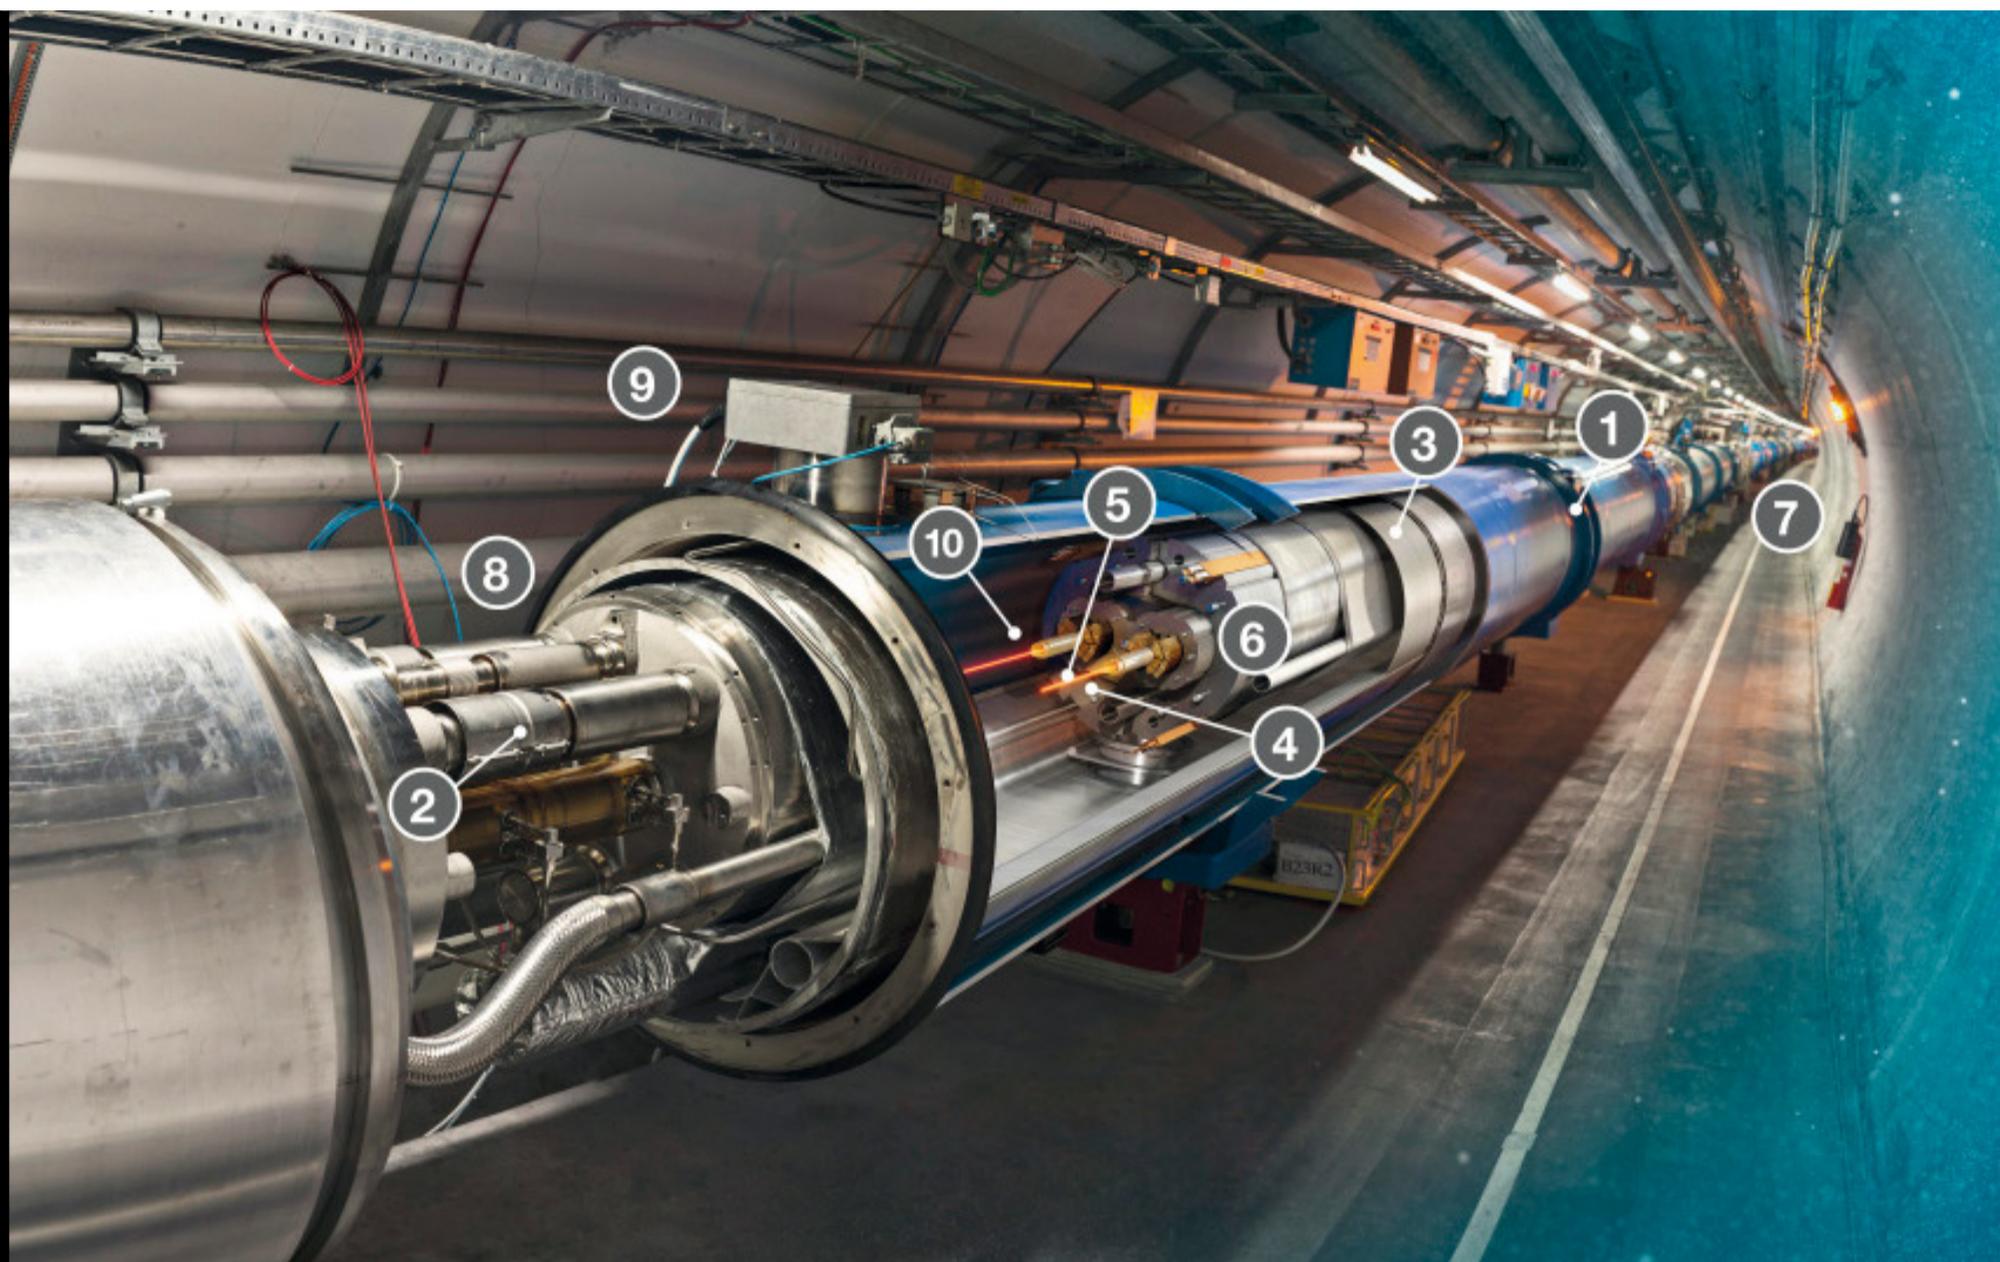
ROOT, THE DATA ANALYSIS FRAMEWORK

This framework now has the following features:

- Scalable, efficient, format independent
- Orthogonal object model
- Object serialization
- Auto Evolution
- Object versioning
- Integrated data compression
- Granularity and easily tunable grouping
- Remote access
- HTTP, HDFS, Amazon S3, CloudFront and Google Storage compliant
- self-describing file format
- ROOT I/O is used to store all the LHC data

CONCLUSION

Since its inception CERN has a long tradition in the deployment systems of storage on a large scale for the scientific community worldwide as a whole. During the first period of operation of the LHC, CERN spent 100 mark Peta-bytes and contributed to the rapid confirmation of the Higgs boson and many other results resulting from the operation of the LHC. Through the first season, the deployment, storage and data management models have evolved significantly and CERN's infrastructure has been evolving, being constantly modernized and optimized to continue meeting new needs as they arise.



LHC SEASON 2: A STRONGER MACHINE

In early 2013, after three years of running, the Large Hadron Collider (LHC) shut down for planned maintenance. Hundreds of engineers and technicians spent two years repairing and strengthening the accelerator in preparation for running at higher energy. Now the world's largest and most powerful particle collider is ready to start up again. So what's new?

1 NEW MAGNETS

Of the LHC's 1232 superconducting dipole magnets, which steer particle beams around the accelerator, 18 have been replaced owing to wear and tear.

2 STRONGER CONNECTIONS

More than 10,000 electrical interconnections between dipole magnets in the LHC have been fitted with shunts – pieces of metal that act as an alternative path for the 11,000 amp current, saving the interconnection if there is a fault.

3 SAFER MAGNETS

The LHC's superconducting magnets have an improved quench protection system. Superconducting magnets conduct electricity without losing energy to resistance, and so can achieve higher magnetic fields. In a quench, a magnet reverts back to a resistive state, releasing a large amount of energy. The quench-protection system in the LHC serves to dissipate this energy in a more controlled manner if it finds any abnormal voltage developing across a magnet.

RADIOFREQUENCY CAVITIES, WHICH GIVE PARTICLES LITTLE KICKS OF ENERGY AS THEY PASS, WILL OPERATE AT HIGHER VOLTAGES TO GIVE THE BEAMS HIGHER ENERGIES.

4 HIGHER ENERGY BEAMS

The energy of collisions in the LHC in 2015 will be 13 TeV (or 6.5 TeV per beam) compared to 8 TeV (4 TeV per beam) in 2012. Higher energy allows physicists to extend the search for new particles and to check previously untestable theories.

5 NARROWER BEAMS

Because transverse beam size – the width of the beam – decreases with increasing energy, beams in the LHC will be more tightly focused, which means more interactions and collisions for the experiments to study.

6 SMALLER BUT CLOSER PROTON PACKETS

There will be fewer protons per packet – or “bunch”: 1.2×10^{11} compared to 1.7×10^{11} in 2012. When dozens of collisions occur at once, it becomes harder for a detector’s computers to disentangle which particle comes

from which collision. With fewer protons in each collision, this problem of “pileup” will be less severe. However, the bunches of protons will be separated in time by 25 nanoseconds compared to 50 nanoseconds. The LHC will thus deliver more particles per unit time as well as more collisions to the experiments.

7 HIGHER VOLTAGE

Radiofrequency cavities, which give particles little kicks of energy as they pass, will operate at higher voltages to give the beams higher energies.

8 SUPERIOR CRYOGENICS

The dipole magnets on the LHC must be kept at low temperature to be in their superconducting state. The cryogenics system has been fully consolidated, with complete maintenance of the cold compressors, as well as an upgrade of the control systems and renovation of the cooling plant.

9 RADIATION-RESISTANT ELECTRONICS

A full maintenance and upgrade of the electrical systems on the LHC included more than 400,000 electrical tests, and the addition of newer, more radiation-tolerant systems.

10 MORE SECURE VACUUM

The inside of the beam pipe is kept under vacuum so that the beam does not crash into molecules in its path. But charged beams can rip electrons from the inside surfaces of the pipe, forming an “electron cloud” that interferes with the beam. To dampen this effect the inside of the beam pipe has been coated with non-evaporable getter (NEG), a material that takes up the electrons. In places, solenoids have been wrapped around the beam pipe to keep electrons from deviating from the sides.



THE SAFETY OF THE LHC

The Large Hadron Collider (LHC) can achieve an energy that no other particle accelerators have reached before, but Nature routinely produces higher energies in cosmic-ray collisions. Concerns about the safety of whatever may be created in such high-energy particle collisions have been addressed for many years. In the light of new experimental data and theoretical understanding, the LHC Safety Assessment Group (LSAG) has updated a review of the analysis made in 2003 by the LHC Safety Study Group, a group of independent scientists. LSAG reaffirms and extends the conclusions of the 2003 report that LHC collisions present no danger and that there are no reasons for concern. Whatever the LHC will do, Nature has already done many times over during the lifetime of the Earth and other astronomical bodies. The LSAG report has been reviewed and endorsed by CERN's Scientific Policy Committee, a group of external scientists that advises CERN's governing body, its Council. The following summarizes the main arguments given in the LSAG report. Anyone interested in more details is encouraged to consult it directly, and the technical scientific papers to which it refers.

COSMIC RAYS

The LHC, like other particle accelerators, recreates the natural phenomena of cosmic rays under controlled laboratory conditions, enabling them to be studied in more detail. Cosmic rays are particles produced in outer space, some of which are accelerated to energies far exceeding those of the LHC. The energy and the rate at which they reach the Earth's atmosphere have been measured in experiments for some 70 years. Over the past billions of years, Nature has already generated on Earth as many collisions as about a million LHC experiments – and the planet still exists. Astronomers observe an enormous number of larger astronomical bodies throughout the Universe, all of which are also struck by cosmic rays. The Universe as a whole conducts more than 10 million million LHC-like experiments per second. The possibility of any dangerous consequences contradicts what astronomers see - stars and galaxies still exist.

MICROSCOPIC BLACK HOLES

Nature forms black holes when certain stars, much larger than our Sun, collapse on themselves at the end of

their lives. They concentrate a very large amount of matter in a very small space. Speculations about microscopic black holes at the LHC refer to particles produced in the collisions of pairs of protons, each of which has an energy comparable to that of a mosquito in flight. Astronomical black holes are much heavier than anything that could be produced at the LHC.

According to the well-established properties of gravity, described by Einstein's relativity, it is impossible for microscopic black holes to be produced at the LHC. There are, however, some speculative theories that predict the production of such particles at the LHC. All these theories predict that these particles would disintegrate immediately. Black holes, therefore, would have no time to start accreting matter and to cause macroscopic effects.

Although theory predicts that microscopic black holes decay rapidly, even hypothetical stable black holes can be shown to be harmless by studying the consequences of their production by cosmic rays. Whilst collisions at the LHC differ from cosmic-ray collisions with astronomical bodies like the Earth in that new particles produced in LHC collisions tend to move more slowly than those produced by cosmic rays, one can still demonstrate their safety. The specific reasons for this depend whether the black holes are electrically charged, or neutral. Many stable black holes would be expected to be electrically charged, since they are created by charged particles. In this case they would interact with ordinary matter and be stopped while traversing the Earth or Sun, whether produced by cosmic rays or the LHC. The fact that the Earth and Sun are still here rules out the possibility that cosmic rays or the LHC could produce dangerous charged microscopic black holes. If stable microscopic black holes had no electric charge, their interactions with the Earth would be very weak. Those produced by cosmic rays would pass harmlessly through the Earth into space, whereas those produced by the LHC could remain on Earth. However, there are much larger and denser astronomical bodies than the Earth in the Universe. Black holes produced in cosmic-ray collisions with bodies such as neutron stars and white dwarf stars would be brought to rest. The continued existence of such dense bodies, as well as the Earth, rules out the possibility of the LHC producing any dangerous black holes.



STRANGELETS

Strangelet is the term given to a hypothetical microscopic lump of 'strange matter' containing almost equal numbers of particles called up, down and strange quarks. According to most theoretical work, strangelets should change to ordinary matter within a thousand-millionth of a second. But could strangelets coalesce with ordinary matter and change it to strange matter? This question was first raised before the startup of the Relativistic Heavy Ion Collider, RHIC, in 2000 in the United States. A study at the time showed that there was no cause for concern, and RHIC has now run for eight years, searching for strangelets without detecting any. At times, the LHC will run with beams of heavy nuclei, just as RHIC does. The LHC's beams will have more energy than RHIC, but this makes it even less likely that strangelets could form. It is difficult for strange matter to stick together in the high temperatures produced by such colliders, rather as ice does not form in hot water. In addition, quarks will be more dilute at the LHC than at RHIC, making it more difficult to assemble strange matter. Strangelet production at the LHC is therefore less likely than at RHIC, and experience there has already validated the arguments that strangelets cannot be produced.

The analysis of the first LHC data from heavy ion collisions has now confirmed the key ingredients used in the LSAG report to evaluate the upper limit on the production of hypothetical strangelets. For more details see this addendum to the LSAG report: Implications of LHC heavy ion data for multi-strange baryon production (2011)

VACUUM BUBBLES

There have been speculations that the Universe is not in its most stable configuration, and that perturbations caused by the LHC could tip it into a more stable state, called a vacuum bubble, in which we could not exist. If the LHC could do this, then so could cosmic-ray collisions. Since such vacuum bubbles have not been produced anywhere in the visible Universe, they will not be made by the LHC.

MAGNETIC MONOPOLES

Magnetic monopoles are hypothetical particles with a single magnetic charge, either a north pole or a south pole. Some speculative theories suggest that, if they do exist, magnetic monopoles could cause protons to decay. These theories also say that such monopoles would

be too heavy to be produced at the LHC. Nevertheless, if the magnetic monopoles were light enough to appear at the LHC, cosmic rays striking the Earth's atmosphere would already be making them, and the Earth would very effectively stop and trap them. The continued existence of the Earth and other astronomical bodies therefore rules out dangerous proton-eating magnetic monopoles light enough to be produced at the LHC.

OTHER ASPECTS OF LHC SAFETY

Concern has recently been expressed that a 'runaway fusion reaction' might be created in the LHC carbon beam dump. The safety of the LHC beam dump had previously been reviewed by the relevant regulatory authorities of the CERN host states, France and Switzerland. The specific concerns expressed more recently have been addressed in a technical memorandum by Assmann et al. As they point out, fusion reactions can be maintained only in material compressed by some external pressure, such as that provided by gravity inside a star, a fission explosion in a thermonuclear device, a magnetic field in a Tokamak, or by continuing isotropic laser or particle beams in the case of inertial fusion. In the case of the LHC beam dump, it is struck once by the beam coming from a single direction. There is no countervailing pressure, so the dump material is not compressed, and no fusion is possible.

FUSION FEAR IMPOSSIBILITY

Concern has also been expressed that a 'runaway fusion reaction' might be created in a nitrogen tank inside the LHC tunnel. There are no such nitrogen tanks. Moreover, the arguments in the previous paragraph prove that no fusion would be possible even if there were.

Finally, concern has also been expressed that the LHC beam might somehow trigger a 'Bose-Nova' in the liquid helium used to cool the LHC magnets. A study by Fairbairn and McElrath has clearly shown there is no possibility of the LHC beam triggering a fusion reaction in helium. We recall that 'Bose-Novae' are known to be related to chemical reactions that release an infinitesimal amount of energy by nuclear standards. We also recall that helium is one of the most stable elements known, and that liquid helium has been used in many previous particle accelerators without mishap. The facts that helium is chemically inert and has no nuclear spin imply that no 'Bose-Nova' can be triggered in the superfluid helium used in the LHC.

NEXT ISSUE



#11
FEBRUARY
1st, 2016

VIEWPOINT
SECURITY IN
INDUSTRIAL
ENVIRONMENTS

COVER STORY
HPC PREDICTIONS
FOR 2016 : CHARTING
THE FUTURE

HOW TO
SETTING UP AN
HPC CLUSTER

TECH ZONE
MAKING A SUCCESS OF AN INDUSTRIAL PROJECT

The winner of our OCZ contest for a 960GB OCZ Saber 1000 SSD Enterprise Drive is Lê Thiên-Hiêp, Scientist at ONERA Aerospace. Congratulations !



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

CONTACTS
editorial@hpcreview.com
subscriptions@hpcreview.com
sales@hpcreview.com

PUBLISHER
Frédéric Milliot

EDITOR IN CHIEF
Joscelyn Flores

EXECUTIVE EDITOR
Ramon Lafleur

CONTRIBUTORS
TO THIS ISSUE
Amaury de Cizancourt
Steve Conway

Roberto Donigni
Ramon Lafleur
Stacie Loving
Solange Belkhayat-Fuchs
Marc Paris

ART DIRECTOR
Bertrand Grousset

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