

CERN's data stores soar to 530M gigabytes

CERN's data requirements are growing at massive rates: Every experiment with its Large Hadron Collider creates 3GB of data per second.

Lucas Mearian (Computerworld (US)) on 14 August, 2015 23:36

 **0 Comments** |  **1408**  **0**  **0**   

Since restarting in June after a two-year upgrade, CERN's Large Hadron Collider (LHC) has been recording about 3GB of data per second, or about 25 petabytes -- that's 25 million gigabytes -- of data per year.

Every time the LHC smashes particles together at near the speed of light in its 16 mile-long chamber, the shattered particles fly off in myriad of directions. Those particlaes leave behind traces in space, like footsteps in snow, which are recorded and later analyzed in a search for the most basic element of matter.

But unlike a camera, which absorbs light in order to produce a photo, the traces that result from particle collisions pass through the LHC's "detectors," leaving many points of interaction in their path. Every point represents an action at a point in time that can help pinpoint the particle's characteristics.

The detectors that record particle collisions have 100 million read-out channels and take 14 million pictures per second. It's akin to saving 14 million selfies with every tick of a watch's second hand.

Needles and haystacks

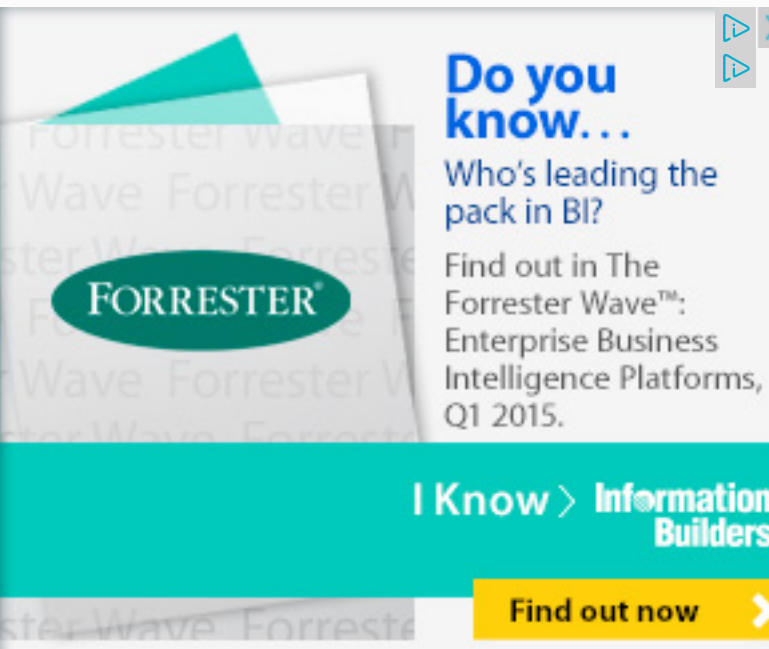
Guenther Dissertori, a professor of particle physics at CERN and the Swiss Federal Institute of Technology in Zurich, said the task of finding matter's most basic particle is vastly more difficult than finding that proverbial needle in a haystack.

"The search for the particle is more than a search for a needle in a haystack. We get 14 million haystacks per second - and unfortunately the needle also looks like hay," Dissertori said. "The amount of data produced at CERN was impressive 10 years ago, but is not as impressive as what's produced today."

Dissertori said CERN's public-private partnerships could solve the expected technological hurdles, including the need for new storage technologies that can save exabytes of data in the future.

Unlike Google or Amazon, two Internet companies that spend billions of dollars every year to develop new technology, CERN has limited money; it's funded by 21 member states and has an annual budget of around \$1.2 billion.

"We have to be very creative to find solutions, Dissertori said. "We're forced to find the best possible ways to collaborate with [the IT] industry and get most out of it."





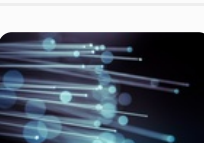


Do you know...
Who's leading the pack in BI?
Find out in The Forrester Wave™: Enterprise Business Intelligence Platforms, Q1 2015.


I Know > Information Builders

Find out now >






EDITOR'S RECOMMENDATIONS

-  Beehives get Intel inside
-  Telstra Health scores hospital contracts in Malaysia, Thailand
-  Optus locks down admin rights
-  NBN still not sure about the state of Telstra's copper
-  NBN 2.0 costs more than expected

READ MORE

-  Symantec Agile Data Centre

FEATURED WHITEPAPERS

-  Tooling up for collaboration
-  7 Benefits you realise with a holistic data protection approach
-  Key reasons to use Software-defined storage and how to get started
-  HOYTS Group entertains more than AU\$300,000 annual cost savings with ShoreTel
-  Sharp demand for managing IT complexity

Almost since its founding, CERN has been developing ways to improve data storage, cloud-technologies, data analytics and data security in support of its research. Its technological advancements have resulted in a number of successful research spin-offs from its primary particle work, including the World Wide Web, hypertext language for linking online documents and grid computing.

Data#3

BRANDPOST

Technical considerations for mobility in the enterprise

More from Data#3 »

Its invention of grid computing technology, known as the [Worldwide LHC Computing Grid](#), has allowed it to distribute data to 170 data centers in 42 countries in order to serve more than 10,000 researchers connected to CERN.

Storing data, sharing data

During the LHC's development phase 15 years ago, CERN knew that the storage technology required to handle the petabytes of data it would create didn't exist. And researchers couldn't keep storing data within the walls of their Geneva laboratories, which already house an impressive 160PB of data.

CERN also needed to share its massive data in a distributed fashion, both for speed of access as well as the lack of onsite storage.

As it has the past, CERN developed the storage and networking technology itself, [launching the OpenLab in 2001](#) to do just that. OpenLab is an open source, public-private partnership between CERN and leading educational institutions and information and communication technology companies, such as Hewlett-Packard and Nexenta, a maker of software-defined storage.

OpenLab itself is a software-defined data center that started phase five of its development cycle this year. That phase will continue through 2017 and tackle the most critical needs of IT infrastructures, including data acquisition, computing platforms, data storage architectures, compute provisioning and management, networks and communication, and data analytics.

A growing grid

In all, the LHC Computing Grid has 132,992 physical CPUs, 553,611 logical CPUs, 300PB of online disk storage and 230PB of nearline (magnetic tape) storage. It's a staggering amount of processing capacity and data storage that relies on having no single point of failure.

In the next 10 to 20 years, data will grow immensely because the intensity of accelerator will be ramped up, according to Dissertori.


"The electronics will be improved so we can write out more data packages per second than we do now," Dissertori said.

Every LHC experiment at the moment writes data on a magnetic tape at the order of 500 data packets per second; each packet is a few megabytes in size. But CERN is striving to keep as much data as possible on disc, or online storage, so that researchers have instant access to it for their own experiments.


"One interesting development is to see how can we implement it with data analysis within our cloud computing paradigm. For now, tests are ongoing on our cloud," Dissertori said. "I could very well imagine in near term future more things done in that direction."

Join the Computerworld newsletter!

Persönlicher Experte für Ihren Mac



JETZT DOWNLOADEN



LATEST JOBS

IDG TechJobs

C	Technical Business Analyst	NSW
C	Business Analyst - Clinical	NSW
F	DevOps Consultant - Microsoft ...	VIC
F	Senior Security Engineer	NSW
F	Information Systems Senior Systems	NZ




Show all jobs

Computerworld newsletter


Join the most dedicated community for IT managers, leaders and professionals in Australia

Enter your email address


Subscribe




RELATED WHITEPAPERS




Shifting gears in the automotive industry



Hot Topic: customer loyalty success in less than a year



5 best practices to optimise sales revenue in 2015



5 things field workers hate and how to fix them

Email address

Join

Tags storage CERN cloud storage

More about CERN Geneva Google Nexenta Technology

Read next




Vodafone CEO repeats telco's call for USO overhaul



Fighting fraud with graph analytics



Kaspersky allegations to breed mistrust
CSO Online



In Pictures: Ultimate guide to Raspberry Pi operating systems, part 1



In Pictures: User guide to Windows 10



The Future Ready IT Series

Jetzt neu in der
gonserBOX:
Kleine Dinge zum
kleinen Preis!

TIEFSTPREIS
GARANTIE

gonserBOX

gonser

Market Place			
	<div>Gartner Symposium/ITxpo – The World’s Most Important Gathering of CIOs and Senior IT Executives 26-29 October 2015 on the Gold Coast, Australia</div>		<div>Using customer personas while maintaining excellence: The key to omnichannel</div>
	<div>Only Citrix securely delivers everything you need for a truly mobile workspace wherever you go</div>		<div>Check out the Log Rhythm's NEW Resource centre New content, infographics, white papers and research</div>

CSO	CIO	ARN	TechWorld	CMO
<div><div>1. Ashley Madison hauled to court in class action suits over data breach</div><div>2. Security Watch: Fortinet launches new WLAN solutions for secure cloud-managed enterprise Wi-Fi</div><div>3. Security Watch: Aronsen Security Group appoints new CSO</div><div>4. Ashley Madison posts \$500k reward for hacker info, police call for hacker help</div><div>5. Startup takes heat over online</div></div>	<div><div>1. How to work with millennials</div><div>2. 5 things you need to know about Windows 10</div><div>3. Maintaining a hybrid cloud environment</div><div>4. 3 reasons to be wary of the Internet of Things</div><div>5. How to do an effective job handover</div></div>	<div><div>1. Macquarie Telecom to offer Riverbed SteelHead SaaS solution</div><div>2. 5 for 15: Adam Smith, F-Secure A/NZ country manager</div><div>3. UPDATED: Avnet scores A/NZ distribution agreement with EMC</div><div>4. IBM dresses up Melbourne Spring Fashion Week 2015</div><div>5. NetDocuments launches local datacentres with NextDC</div></div>	<div><div>1. Telstra Health scores hospital contracts in Malaysia, Thailand</div><div>2. How 'low code' can drive innovation</div><div>3. Alibaba adds artificial intelligence capability to its cloud offerings</div><div>4. Ashley Madison hauled to court in class action suits over data breach</div><div>5. Beehives get Intel inside</div></div>	<div><div>1. CMO Interview: Why Hootsuite's marketing and product teams are inseparable</div><div>2. Myer appoints Mark Cripsey as new chief digital and data officer</div><div>3. Qantas announces further investment in customer engagement</div><div>4. AdRoll retargeting integrates with Shopify</div><div>5. Oracle acquires Maxymiser to increase digital engagement and revenue</div></div>

Join the
Computerworld
Community

Sign up



- | | | |
|---|-----------------------------|-----------------------------------|
| Editorial Contacts | Events | IT Media Releases |
| Advertising Information | Whitepapers | Slideshows |
| Privacy Policy | News | Videos |
| RSS | Webinars | |
| Newsletters | Zones | |



Copyright 2015 IDG Communications. ABN 14 001 592 650. All rights reserved.
Reproduction in whole or in part in any form or medium without express written permission of IDG Communications is prohibited.
IDG Sites: [PC World](#) | [GoodGearGuide](#) | [CIO](#) | [CMO](#) | [CSO](#) | [Techworld](#) | [ARN](#) | [CIO Executive Council](#) | [IDG Education](#) | [IDG Government](#) | [IDG Health](#)

