CERN
Clouds & particles

The European Organization for Nuclear Research (CERN) is the center of the particle physics universe (so to speak), and has been in the news a lot lately due to its recent discovery of a particle consistent with the long-sought Higgs boson, dubbed by some the "God particle." However, as is often the case in physics, the surmounting of one obstacle only reveals another. Scientists will now attempt to use this discovery to help unravel some fundamental forces of existence.

By Xue Hua & Hu Yuhai **Editor: Xu Shenglan** xushenglan@huawei.com

> In 2007, CERN's research data totaled seven petabytes. In 2009, it reached fifteen petabytes, and would reach 25 petabytes three years later.

CERN, headquartered in Switzerland, is the world's largest particle physics research organization, boasting over half of the scientists in the international particle physics community as its members. It has been the cradle of numerous scientific breakthroughs since its founding in 1954, including the world's first proton-proton collision in 1971, the discovery of neutral currents in 1973, and the finding of the W and Z particles in 1983, to name just a few. On July 4th, 2012, the laboratory declared that two general-purpose experiments conducted on A Toroidal Large Hadron Collider (LHC) Apparatus (ATLAS) and the Compact Muon Solenoid (CMS) had evinced the presence of a new subatomic particle with a mass of roughly 125-126 gigaelectronvolts (GeV) and the expected profile of the long-sought Higgs boson, which is thought to give mass to all matter; a discovery that led Stephen Hawking, the preeminent physicist and only authentic celebrity in the hard sciences, being out a hundred dollars, as the iconic author of *A Brief History of Time* had previously doubted the existence of the Higgs boson.

CERN's bottleneck in storage

However, this great discovery is only one of many that CERN hopes to achieve. To accommodate the needs of its research in high-energy particle physics, the scientific community built a very large-scale computing grid system, tasked with nothing less than helping to unlock the secrets of the universe.

CERN's data center, also known as the Worldwide LHC Computing Grid (WLCG) Tier-0, forms the core of the organization's global computing resources, and presently enables the storage and analysis of more than 20 petabytes of LHC data per year.

In collaboration with the globe's ICT leaders, CERN has established CERN openlab, whose mission is to accelerate the development of cutting-edge IT solutions for use by the global particle physics community. In recent years, a tremendous surge in data at CERN has severely challenged the scalability and reliability of its storage systems. In 2011, CERN openlab began to evaluate new storage technologies using the following criteria.

First, a highly-reliable storage system was needed to store research data accumulated over years of experimentation, including invaluable raw data from over 16 million particle collisions. Second, said system must be highly scalable to accommodate the data that is to come. In 2007, CERN's research data totaled seven petabytes. In 2009, it reached fifteen petabytes, and would reach 25 petabytes three years later. Third, the storage system must exhibit exceptional throughput and read/write performance, so that retrieval is fast, even in light of the vast scale of data being analyzed. And lastly, it must be both seamless and secure, as data may be stored in different databases at different institutions in different countries.

Huawei UDS cloud storage

CERN openlab's fourth phase, which started in January 2012, heralded the launching of more in-depth cooperative activities in projects involving automation & control, databases, networking, and platforms, for all of which the storage system is indispensable.

As a storage expert, Huawei is committed to

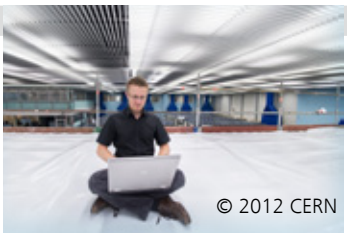

developing innovative next-generation cloud solutions. By joining CERN openlab as a contributor, Huawei has developed a cloud-based Universal Distributed Storage (UDS) system to meet the challenges of exascale data storage (one exabyte equals 1,048,576 terabytes).

In addition to this scalability, Huawei UDS cloud storage features tremendous reliability and broad compatibility. It also employs dense nodular distribution and is based on the energy-efficient Advanced RISC Machine (ARM) architecture. UDS enables storage and sharing of big data on its mass object-based storage infrastructure, which integrates object-based storage, a P2P distributed storage engine, and clustering applications. The system also intelligently adjusts the workload at each node to break through previous performance bottlenecks, so system performance grows evenly with each capacity expansion.

Huawei UDS cloud storage system features reliability to the tune of 99.999999999% (that's eleven nines), while facilitating redundancy across data centers. It also ensures data security throughout the system's lifecycle through multi-layer protection at the node, storage, and data center level through multiple copies, erasure codes, and multi-layer encryption verification involving the interface, storage, and transmission processes.

This solution is also supremely cost-effective. Maintenance costs are curtailed through the use of desktop hard drives, intelligent & automatic fault recovery, energy-efficient ARM chips, efficacious cooling technology, and various other self-maintenance measures (which effectively enhance storage availability & reliability), while compatibility with S3 interfaces, NFS/CIFS protocols, and mainstream backup software ensure an E2E public cloud solution that is easy to use, while meeting the storage access requirements of different applications.

UDS is ready for greater challenges

In early 2012, Huawei delivered its UDS cloud storage system to CERN, with installation and benchmark performance evaluations completed three months later. The read/write performance proved excellent in large-scale data environments, while scalability proved not too shabby either. On the whole, these scores were deemed worthy of the very organization that gave birth to the World Wide Web.

"CERN is hitting the technology limits for resource-intensive simulations and analysis. Our collaboration with Huawei shows an exciting new approach, where their novel architecture extends the capabilities in preparation for the exascale data rates and volumes we expect in the future," said Bob Jones, Head of CERN openlab.

"Establishing the link with CERN openlab gave us a fantastic opportunity to further develop our cloud storage products, and proved their worth in the extreme scientific research and mass data environment," added James Hughes, Huawei's Chief Architect of Cloud Storage.

CERN is pushing further into the strange boundaries of the quantum realm. It is now carrying out more in-depth studies to better understand the aforementioned new particle so that other enigmas of the universe do not remain so. This should give Huawei's UDS cloud storage system further opportunities to distinguish itself, on a scale once unimagined.