

IBM dévoile sa technologie Software Defined Storage pour l'ère du Big Data

Paris, France - 13 mai 2014: IBM dévoile aujourd'hui une nouvelle gamme de produits pour le Software Defined Storage (SDS) qui réduisent les coûts de fonctionnement tout en permettant aux organisations d'accéder et de traiter tout type de données, sur tout type de périphérique de stockage, partout dans le monde. En effet, IBM a développé une solution nommée "Elastic Storage" qui permet de réduire les coûts de stockage des datacenters jusqu'à 90% en transférant automatiquement les données vers l'unité de stockage la plus économique.

Cette innovation, qui a par ailleurs alimenté IBM Watson ainsi que les plus puissants supercalculateurs, est désormais disponible à la commercialisation.

IBM Unveils Software Defined Storage Technology for Era of Big Data

- *Breakthrough storage software provides infinite scaling across all data types*
- *Changes economics of the datacenter by reducing storage costs up to 90 percent*
- *Innovation that powered IBM Watson and top supercomputers now available commercially*

BOSTON - 13 May 2014: IBM (NYSE: [IBM](#)) today unveiled a portfolio of software defined storage products that deliver improved economics at the same time they enable organizations to access and process any type of data, on any type of storage device, anywhere in the world.



One technology in the portfolio, codenamed "Elastic Storage," offers unprecedented performance, infinite scale, and is capable of reducing storage costs up to 90 percent by automatically moving data onto the most

economical storage device.

Born in IBM Research Labs, this new, [patented](#) breakthrough technology allows enterprises to exploit – not just manage – the exploding growth of data in a variety of forms generated by countless devices, sensors, business processes, and social networks. The new storage software is ideally suited for the most data-intensive applications, which require high-speed access to massive volumes of information – from seismic data processing, risk management and financial analysis, weather modeling, and scientific research, to determining the next best action in real-time retail situations.

*"Digital information is growing at such a rapid rate and in such dramatic volumes that traditional storage systems used to house and manage it will eventually run out of runway," said **Tom Rosamilia, Senior Vice President, IBM Systems and Technology Group.** "Our technology offers the advances in speed, scalability and cost savings that clients require to operate in a world where data is the basis of competitive advantage."*

Software-defined storage is a set of software capabilities that automatically manage data locally and globally, providing breakthrough speed in data access, easier administration and the ability to scale technology infrastructures quickly and more cost-effectively as data volumes expand. In addition, these advances can work with any company's storage systems to provide automated and virtualized storage.

Game-Changing Technology

The foundations of today's Elastic Storage were used for the *Jeopardy!* television match between IBM's [Watson](#) and two former *Jeopardy!* champions. For the show, IBM's Watson had access to 200 million pages of structured and unstructured data, including the full text of Wikipedia. By using Elastic Storage capabilities, around five terabytes of Watson's "knowledge" (or 200 million pages of data) were loaded in only minutes into the computer's memory.

A key reason these capabilities were chosen for the Watson system that competed on *Jeopardy!* was its scalability, the architectural limits for which stretch into the thousands of "yottabytes." A yottabyte is one billion petabytes, or the equivalent of a data center the size of one million city blocks, which would fill the states of Delaware and Rhode Island combined.

IBM Research has demonstrated that Elastic Storage can successfully scan 10 billion files on a single cluster in just 43 minutes – a technology demonstration that translates into unequalled performance for clients analyzing massive data repositories to extract business insights.

At its core, Elastic Storage builds on IBM's global file system software to provide online storage management, scalable access, and integrated data governance tools capable of managing vast amounts of data and billions of files. For example, Elastic Storage also exploits server-side Flash for up to six times increase in performance than with standard SAS disks. This feature recognizes when a server has Flash storage and automatically uses that Flash as cache memory to improve performance.

Elastic Storage virtualizes the storage allowing multiple systems and applications to share common pools of storage. This enables transparent global access to data without the need to modify applications and without the need for additional and often disruptive storage management applications. Since Elastic Storage is not reliant on centralized management to determine file location and placement, customers can have continuous and highly-available access to data in the event of software or hardware failures.

For the National Center for Atmospheric Research's Computational and Information Services Laboratory (CISL), growing data volumes are part of its DNA. The organization, which stores and manages more than 50 petabytes of information between its Wyoming and Colorado centers, relies on Elastic Storage to give researchers fast access to vast amounts of diverse data.

*"We provide computational, educational, and research data services for geosciences to more than 1,000 users at more than 200 different sites," said **Pamela Gillman, manager, Data Analysis Services Group, CISL**. "The IBM global file system software has enabled scalable, reliable and fast access to this information. That has dramatically improved the performance of the different functions, as well as the organization as a whole."*

A key component of Elastic Storage is its ability to automatically and intelligently move data to the most strategic and economic storage system available. Through policy-driven features and real-time analytics, for example, Elastic Storage can automatically move infrequently-used data to less expensive low-cost tape drives, while storing more frequently-accessed data on high-speed [Flash systems](#) for quicker access. Such policy-driven features can provide cost savings of up to 90 percent.

In addition, the software features native encryption and secure erase, which ensures that data is irretrievable to comply with regulations such as HIPAA and Sarbanes-Oxley.

Through its support of OpenStack cloud management software, Elastic Storage also enables customers to store, manage and access data across private, public and [hybrid clouds](#) for global data sharing and collaboration. In addition to supporting OpenStack Cinder and Swift access, Elastic Storage supports other open APIs such as POSIX and Hadoop.

While traditional storage systems must move data to separate designated systems for transaction processing and analytics, Elastic Storage can automatically balance resources to support both types of application workloads, including Hadoop based analytics. This dramatically speeds analysis and eliminates the costly and time-consuming process of producing duplicate copies of data.

Elastic Storage software will also be available as an IBM SoftLayer cloud service later this year.

About IBM Big Data and Analytics

Each day, we create 2.5 quintillion bytes of data generated by a variety of sources -- from climate information, to posts on social media sites, and from purchase transaction records to healthcare medical images. At IBM, we believe that data is emerging as the world's newest resource for competitive advantage, and analytics is the key to make sense of it. IBM is helping clients harness [Big Data](#) and [analytics](#) to provide insights needed to make better decisions, create value, and deliver that value to clients and society. IBM has the world's deepest and broadest portfolio of Big Data and analytics technologies and solutions, spanning services, software, research and hardware. For more information about IBM Big Data and analytics, visit <http://ibm.co/bigdataanalytics>. Follow IBM Big Data and analytics on Twitter @IBMbigdata and @IBMAalytics.
