


SOLUTION BRIEF

Intel® Distribution for Apache Hadoop* software
Intel® Manager for Apache Hadoop software
Intel® Xeon® Processor E5 Family
Cloud Computing Solutions



Pioneers in Big Data

Intel® Distribution for Apache Hadoop* software enables Cloudwatt to develop big data Hadoop-enhanced Infrastructure-as-a-service offering running on OpenStack



"Intel® Manager for Apache Hadoop* software allows us to share the Hadoop cluster between multiple users while maintaining guaranteed SLAs."

*Fayçal Boujemaa,
Research and Development Manager,
Cloudwatt*

Challenges

- **Big Data.** Cloudwatt wanted to develop a big data platform-as-a-service offering to help its customers meet the challenge of operating in what it calls the "Data Tsunami" era
- **Open Source.** It needed to find a hardware and software solution that would enable it to develop this big data offering on an OpenStack cloud operating system (CloudOS)

Solutions

- **Proof of Concept.** Cloudwatt carried out a proof of concept (PoC) on Intel® Distribution for Apache Hadoop* software running on the Intel® Xeon® processor E5 family and Intel® Solid-State Drives (Intel® SSDs) DC S3500 series

Technology Results

- **Faster provisioning.** Using Intel Distribution for Apache Hadoop software together with Intel SSDs, Cloudwatt's internal tests showed the ability to provision a cluster in less than 2 minutes
- **Solid state drives.** When analyzing data and adding Intel SSDs to its architecture, Cloudwatt was able to significantly improve the overall performances of its Intel Distribution for Apache Hadoop software running on OpenStack system, achieving the same results as Intel Distribution for Apache Hadoop software working in bare-metal mode with traditional hard disk drives (HDDs)
- **Guaranteed service.** Intel® Manager for Apache Hadoop software allows Cloudwatt to share the Hadoop cluster between multiple users while maintaining guaranteed service-level agreements (SLAs)

Business Value

- **Broader portfolio.** Intel Distribution for Apache Hadoop software enables Cloudwatt to expand its cloud computing offering beyond storage to big data analytics, attracting new customers and opening up new revenue streams
- **Better contacts.** Cloudwatt will be able to benefit from Intel's rich cloud ecosystem to further develop its cloud offering

Cloudwatt was established in September 2012 by Orange, Thales and the French State (funded by the French Strategic Investment Fund). It offers cloud computing infrastructure to small or large companies, public or private, across France and Europe. It has already launched a storage-as-a-service offering based on Dell servers, Intel Xeon processors and OpenStack. To facilitate its strategic goal of extending its cloud infrastructure offering, Cloudwatt is continuously on the look-out for the best and latest technology.

Fayçal Boujemaa, research and development manager at Cloudwatt, explains: "As you may expect, we have regular meetings with Intel to discuss its latest hardware developments and the benefits they can bring to our business, as well as those of our customers. During these conversations Intel also brought our attention to its software offering. We were impressed by Intel's work in this space

and started to think about the ways in which we could start to take advantage of this.

One of our strategic goals for the coming years was to develop, in conjunction with ISVs, a big data analytics offering based on OpenStack," continues Fayçal. "This would ultimately enable our customers to meet the challenge of doing business in the Data Tsunami era. Intel Distribution for Apache Hadoop software seemed like a natural fit in realizing this strategy."

Intel Distribution for Apache Hadoop software provides distributed processing and data management for enterprise applications that analyze massive amounts of diverse data. It includes Apache Hadoop and other software components along with enhancements and fixes from Intel. With optimizations for Intel Xeon processors, Intel SSD storage, and Intel® 10GbE Network Connection, it offers up to a 30x boost in Hadoop performance.



Cloudwatt will benefit from Intel's rich cloud ecosystem

"There are many different distribution platforms for Apache Hadoop in the marketplace," explains Fayçal, "but from a technical perspective we were particularly attracted to the Open Source platform from Intel. Intel Manager for Apache Hadoop software allows us to share the Hadoop cluster between multiple users while maintaining guaranteed SLAs.

Big data analytics are highly computing-intensive and require many read/write functions," continues Fayçal. "Traditional HDDs could become a bottleneck for customers who want to scale their analytics capability. When one of our customers asks for extra resources, we need to be able to provision this fast. This is critical to the success of our big data offering. We were eager to investigate the benefits of Intel SSDs in helping us to achieve this."

Together with Intel, Cloudwatt carried out a PoC on Intel Distribution for Apache Hadoop software, comparing its performance with HDDs and Intel SSDs DC S3500 series. Using Intel Distribution for Apache Hadoop software together with Intel SSDs, Cloudwatt's internal tests showed that it was able to provision a cluster in less than two minutes. When analyzing data and adding Intel SSDs to its architecture, Cloudwatt was able to significantly improve the overall performances of its Intel Distribution for Apache Hadoop software running on OpenStack system, achieving the same results as Intel Distribution for Apache Hadoop software working in bare-mode metal with traditional HDDs.

In addition to providing the hardware and software, Intel was also on hand throughout the PoC to offer Cloudwatt support and advice. Intel experts from

all over Europe joined forces to ensure Cloudwatt was able to successfully demonstrate its Hadoop infrastructure-as-a-service at the OpenStack Summit in Hong Kong.

Cloudwatt now plans to roll out Hadoop infrastructure-as-a-service over the next six to nine months. The offering will be based on Dell hardware running the Intel Xeon processors E5 family and OpenStack. Customers will be able to choose from different performance levels, the most basic of which will run on HDDs and the fastest of which will run on Intel SSDs DC S3500 series.

Ultimately, Intel Distribution for Apache Hadoop software is enabling Cloudwatt to expand its cloud computing offering beyond storage to big data analytics, attracting new customers and opening up new revenue streams. Cloudwatt is also hoping to benefit from Intel's rich cloud ecosystem to develop this offering even further.

Since many of Cloudwatt's customers are already expressing an interest in protecting data beyond classic security protection offered by the network, Cloudwatt is now investigating the security benefits of Intel® Advanced Encryption Standard Instructions (Intel® AES-NI).


Lack of integrated support for strong data security has been a serious roadblock to implementation for many businesses. Encryption and decryption are compute-intensive processes that traditionally add considerable latency and consume substantial processing resources. Intel Distribution for Apache Hadoop software eliminates much of the latency and greatly reduces the load on the processors using Intel AES-NI built into the Intel Xeon processor.

Lessons Learned

Intel's next-generation data center offering goes beyond traditional hardware to software solutions that help organizations meet the challenge of, among other things, big data. Intel® Distribution for Apache Hadoop software is the only distribution built from silicon up to enable the widest range of analytics on Apache Hadoop. Intel® Manager for Apache Hadoop software allows organizations to share the Hadoop cluster between multiple users while maintaining guaranteed SLAs. Together with Intel® Xeon processors, Intel® SSDs, and Intel® 10GbE Network Connection, Intel Distribution for Apache Hadoop software offers even greater Hadoop performance.

Visit Intel's Technology Provider website at www.inteltechnologyprovider.com

Find the solution that's right for your organization. Contact your Intel representative, visit Intel's Business Success Stories for IT Managers (www.intel.co.uk/itcasestudies) or explore the Intel.co.uk IT Center (www.intel.co.uk/itcenter).



This document and the information given are for the convenience of Intel's customer base and are provided "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed, or contained herein. Intel® products are not intended for use in medical, lifesaving, life-sustaining, critical control, or safety systems, or in nuclear facility applications.

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>.

No computer system can provide absolute security. Requires an enabled Intel® processor and software optimized for use of the technology. Consult your system manufacturer and/or software vendor for more information.

Copyright © 2013, Intel Corporation. All rights reserved. Intel, the Intel logo, Intel Xeon, Xeon inside, Intel Solid-State Drives, Intel Distribution for Apache Hadoop software, Intel Manager for Apache Hadoop software, and Intel Advanced Encryption Standard Instructions are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

1213/JNW/RLC/XX/PDF

329935-001EN