


CASE STUDY

Intel® Xeon® Processors

Intel® Distribution for Apache Hadoop* Software



China Mobile Guangdong Gives Subscribers Real-Time Access to Billing and Call Data Records

Intel® hardware and software solutions help China Mobile Guangdong build a new high-performance, reliable and cost-effective detailed billing statement inquiry system as it gears up for continuous customer growth



China Mobile Group Guangdong Co., Ltd.

is a subsidiary of China Mobile Communications Corporation. It is the leading provincial-level communications operator in terms of revenue.

CHALLENGES

- **Eroding profit margins.** The high cost of maintaining the existing billing system erodes billing business unit profitability.
- **Poor quality of service.** The current high-touch customer support model is not scalable to match explosive growth, risking dissatisfied customers who might switch to the competition.
- **Address scale issues.** Current relational database management system (RDBMS) solutions are failing to deliver storage scalability and real-time queries needed to meet subscribers' needs.

SOLUTION AND BENEFITS

- **Optimize hardware performance for big data.** Replace RISC* platform with Intel® Xeon® processor 5600 series generic compute platform tuned for Apache Hadoop software to lower TCO and increase performance.
- **Near-real-time Hadoop-powered analysis.** Use Intel® Distribution for Apache Hadoop* (Intel® Distribution) to remove data access bottlenecks and uncover customer usage patterns for more targeted marketing promotions.
- **Expand Storage with Hadoop HBase*.** Intel Distribution's Big Table enhancements to Hadoop HBase deliver auto-splitting of data tables across nodes that expand storage at lower cost.

Growing Pains

As in many countries, Chinese communication service providers have experienced explosive growth as consumers embrace mobile devices, fast 3G and 4G connectivity, and the ability to self-service or inquire about account-related information. In parallel, traditional customer support operations continue to grow as a means to deliver personalized customer support. With such development, the challenge for communication service providers is to provide a broad range of information assets and instant analysis

capabilities that deliver more information and targeted offerings for the consumer.

Keeping its clientele in mind, China Mobile Guangdong has always aimed to elevate customer experience by providing services that target their needs. The company cites its mobile operation support system, and customer service in general, as key assets to provide a differentiated offering to its customers. Within this system, the detailed billing system component delivers one of the

Intel® Distribution software enterprise-class solutions and Intel® Xeon® processor 5600 series helped China Mobile Guangdong build a new detailed billing statement inquiry system, which lays a firm foundation for constant improvement of its customer services while increasing its volume of business



“China Mobile Guangdong’s billing statement inquiry system, which is built on Intel® Distribution for Apache Hadoop software optimized for Intel Xeon processors, has delivered the scale and real-time information our operators need to service a growing mobile subscriber base.”

Tang Hui
Project Manager
China Mobile Group Guangdong Co., Ltd.

most important touch points with customers and for the billing support staff.

“The detailed billing system provides quite a broad range of inquiry services for users, including voice calls, Internet surfing, short text messaging and data communications related to online retrieval of monthly billing/CDR information,” said Tang Hui, project manager, China Mobile Guangdong. “In addition, with our business’s expansion, the number of our 3G users and billing data records has increased the volume and velocity of our inquiry data requests.”

High maintenance costs threaten storage expansion

To meet the demands of its growing customer base, China Mobile Guangdong focused on expanding the storage and computational capacity of its billing system. However, new investments and maintenance costs to upgrade the existing RISC minicomputer-based platform of the billing system would surpass allocated budgets. Moreover, the upgrade and expansion process would need the computer to stop, which would interrupt business and affect service quality.

China Mobile Guangdong also evaluated a new architecture based on massively parallel processing (MPP). Traditional data warehouse (Dw) technologies were also evaluated but ruled out for three main reasons:

1. Total cost of ownership (TCO) to scale up storage for this centralized RDBMS approach
2. Need for real-time analytics and decision support
3. Solid disaster recovery and failover

Lack of reliable service or downtime would put China Mobile Guangdong’s credibility at risk. If failure occurs in storing mass data in the billing statement inquiry system, the repair might take several days. If the data were lost, the company’s brand and customer loyalty would suffer.

The final step for China Mobile Guangdong was to evaluate the emerging big data Hadoop open source frameworks where storage and analytical processing are distributed across 1-n node clusters. Early research indicated big data might be a good fit for unstructured big data processing, but the current billing system dealt with large call data record (CDR) files that had static structures. On the other hand, industry TCO modeling between the MPP DW system and Hadoop presented a stark contrast in data processing expense: USD57,000 per terabyte for traditional MPP versus less than USD1,000 for Hadoop big data. As an open source solution, Hadoop solved another requirement, to base future solutions on open, commodity server hardware that could be easily upgraded to take advantage of future lower-cost processors. Avoiding vendor lock-in was a prime directive. Final concerns related to the enterprise readiness of an open source Hadoop distribution, along with the specialized skills and effort required to patch, deploy, tune, and support this effort into the future.

Transition from RISC platform to Hadoop software solutions improves billing system

The China Mobile Guangdong team focused on the server hardware problem first, since the warranty period for the RISC system was fast approaching and high future expansion capacity was a foundational problem to solve. The team found its solution in Intel Xeon processor 5600 series, whose benchmark performance data provides a solid match to replace the RISC-based IBM POWER Series* minicomputer system with high-performance, commodity hardware.

During the engagement, the team learned that Intel had an enterprise software offering, Intel Distribution, which contains the full distribution from the Apache Hadoop open source project and a value-added monitoring

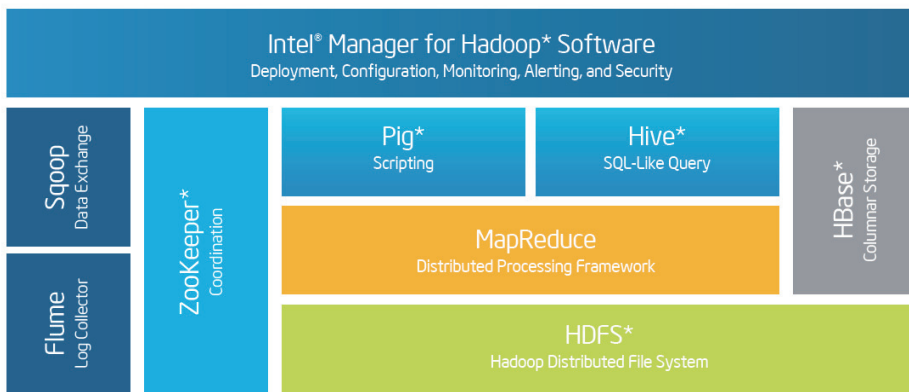


Figure 1. Intel Distribution for Apache Hadoop* software components.

and management console, Intel Manager, for Hadoop. This subscription-based offering was fully supported by Intel and would take the time-intensive process of distributing updates and maintenance out of the picture.

The concerns over team experience were alleviated by engaging Intel's professional service options for proof of concept (POC) design, installation, deployment and configuration. Post deployment the Guangdong team would be trained to be self-sufficient in its operations with Intel's Administrator Training and Developer Training education packages.

While other point solution Hadoop distribution vendors were considered, the fact that Intel had both hardware and software solutions designed to work as a total big data solution, performance-optimized and tuned for common vertical analytical models, became a deciding factor.

China Mobile Guangdong eventually chose Intel Distribution and Intel Xeon processors to build the new system with 133 servers.

China Mobile Guangdong gives subscribers real-time access to billing and call data records

Intel-optimized big data hardware and software solutions enabled China Mobile Guangdong's individual subscribers to access and pay their bills online and retrieve 6 months worth of call data records (CDR) in near-real time. Intel solutions enabled this at a fraction of the price of the incumbent RISC*-based technologies in a way that can economically scale to several hundred million subscribers. Moreover, analytics on cell tower performance, subscriber phone device preferences, and customer usage patterns delivered new customer segmentation and targeted promotions.

Business benefits

Hardware. The increased performance of the solution enabled China Mobile Guangdong to invest in less hardware overall, increased energy efficiency savings, and provided a future upgrade path that was significantly less than the old RISC-based system.

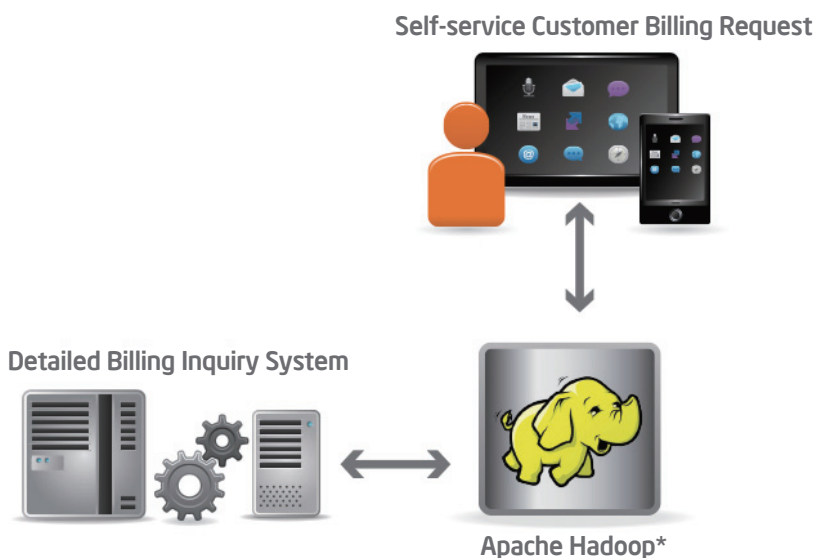


Figure 2. China Mobile end users can now access their accounts, pay bills and retrieve 6 months worth of call data records in real time with Intel-optimized big data hardware and software solutions

Solution Highlights

- **Higher performance.** Based on the Intel Distribution's underlying optimization algorithms, the applications became more efficient and the computing storage was distributed in a more balanced way. The auto-tuning configurations controlled by the Intel Manager for Hadoop installation programs helped optimize performance seamlessly with the Intel Xeon processor.
- **Stable operation.** The fully-tested, enterprise-level Hadoop release has ensured long-term stable operation. Integrated with the latest open source, the consistency among the various components was guaranteed and fully supported by Intel, simplifying operations management.
- **Easy to manage.** Intel Manager for Hadoop delivered a browser-based cluster installation and management controls to solve the problems resulting from difficult management of the open source edition. It provided system alerting for abnormalities in the Web pages, mails, and SMS for proactive customer service.
- **Intel enhancements.** The Intel Distribution contains hundreds of enhancements from Intel, including an HBase database virtual big table which is cross data center. Intel aims to give back to the open source community by focusing on an open big data platform for hardware and software .
- **New information products.** The new analytical capabilities have opened the door to a wide range of new information products that can be sold to targeted customer segments versus delivering a single offering to all customers. Analytics provide the correlation between actual customer usage, preferences, and cost to deliver services by customer segment.

Storage. To increase the expansion capacity, the Intel Distribution's distributed database component (HBase*) stored detailed billing information in a large table of HBase, while the HBase cluster provided detailed data writing and inquiry services. When the entire system needs a capacity expansion, the data nodes can be increased, and the system can automatically match the new server into the whole array on in real time. HBase's large table automatically splits the data table and deploys the data table's access service on all the data nodes of the clusters. Therefore, the new detailed billing statement inquiry system provides a cost effective scale-up capacity expansion for China Mobile Guangdong.

Performance. Even though China Mobile Guangdong is focused on structured data today, Hadoop performed flawlessly and introduced a platform to process unstructured data needs into the future. By improving the mass data processing performance, the Intel Distribution enabled distributed data access from the HBase database engine of different server nodes. Through the high aggregate network bandwidth of the cluster's distribution services, this solution brought about high-speed HBase database access. Moreover, HBase's Share-Nothing framework removed the bottleneck of database file access bandwidth and interlocks. The database access throughput was increased from several to hundreds of times, resulting in extremely high data throughput. The speed at which China Mobile Guangdong's detailed billing statement inquiry system can now retrieve queries is 300,000 records per second, while the insert speed is 800,000 records per second. The system is seamlessly handling 30TB of

subscriber billing data per month, with each table supporting billions of records. The inquiry performance has increased by 30 times, significantly improving the new system's processing performance even as millions of subscribers query the system.

Reliability. Up-time and disaster recovery have dramatically improved. The system does not need to be maintained on a centralized computer or single MPP DW, which provides improved disaster recovery across clusters that average 80 nodes each. The database file of HBase can be reliably stored through Hadoop Distributed File System* (HDFS*). HDFS not only provides high aggregate bandwidth file access through distributed storage/retrieval, it also copies the file's information three times to ensure that it can still provide data and file reading and writing when the hardware and network fail to work. The self-restoration function of HDFS can always make the data's information keep three copies in the clusters for good, so the data will not be lost. At the same time, the database service in each HBase Region Server can transfer to other services in real time when the server fails to work.

Analytics. China Mobile Guangdong can now use large volumes of call data records stored in HBase for analytics to effectively drive business decisions. The business intelligence data is derived from 15 initial analytics queries across network planning and operations, service assurance, billing, marketing, revenue assurance, and customer management. Using a combination of real-time and non-real-time data aggregation, extraction, analytics mining, and reporting, they have uncovered actionable insights from customer behavior. As an example, they can analyze subscriber

usage data combined with demographics to segment customers for personalized service recommendations and target higher-paying customers for new offerings. Similarly, they can proactively monitor network usage to uncover performance bottlenecks and discover which sites and users are incurring the highest data charges.

"Our new billing statement inquiry system has delivered a low TCO, high expansion capability, and high processing performance, thus laying a very solid foundation to enable China Mobile Guangdong to constantly improve its customer services in the context of a high volume of businesses. The successful accomplishment of this new system will undoubtedly bring our service quality to a new level," added Hui.

To find out more on Intel's optimized solutions for big data visit www.intel.com/bigdata

E-mail: asipcustomercare@intel.com

855-229-5580 Toll Free US

+1-425-888-0426 International

Find the solution that's right for your organization.

Contact your Intel representative, visit Intel's

Business Success Stories for IT Managers

(www.intel.com/itcasestudies)

or explore the Intel.com IT Center

(www.intel.com/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 NONINFRINGEMENT 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.

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.

© 2012, Intel Corporation. All rights reserved. Intel, the Intel logo, and Intel Xeon inside are trademarks of Intel Corporation in the U.S. and other countries.

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

0912/EMM/PMG/XXPDF

327967-001EN