

CASE STUDY

Intel® Intelligent Storage Acceleration Library

Intel® VTune™ Amplifier XE

Intel® Ethernet 10 Gigabit Converged Network Adapter

Intelligent Storage



Look Inside.™

Software optimization significantly boosts storage performance

Intel software development tools help Tencent improve storage performance, unleash hardware performance, and maximize return on investment

Tencent 腾讯

“Intel® Xeon® processors, as well as a series of other optimizations, have helped our TFS* accomplish a qualitative leap in its performance, effectively saving a large amount of storage space, considerably reducing storage costs and energy consumption of the data center, and vigorously supporting construction of a green data center. Meanwhile, Intel's optimization program can also help us dramatically reduce development costs with its multi-platform compatibility.”

Li Bo,
Head of the File Storage Team,
Architecture Platform Department,
Tencent

Founded in 1998, Tencent is one of China's largest comprehensive Internet service providers and known for popular products like QQ instant messenger*, Weixin* (and a similar version known as WeChat* overseas), and online games. Over the years, Tencent has kept making enormous investments in storage technology, data mining, and related fields and been committed to providing Internet users with the best possible experience. Behind these products, Tencent File System* (TFS*) serves at the core of file services necessary for many businesses. With hundreds of millions of users, TFS is facing huge pressure in terms of its performance and capacity. Since Tencent Data Center is mainly based on Intel® architecture, Tencent is working with Intel to optimize the performance of TFS.

CHALLENGES

- **Storage pressure and performance bottlenecks.** The traditional triple redundancy solution adopted by Tencent has put tremendous pressure on the storage system. A growing amount of cold data has further decreased its cost/performance.

SOLUTIONS

- **Intel® Xeon® processors and software optimizations.** With powerful processing capacity, Intel Xeon processors enabled Tencent to successfully transform its storage solution to the erasure code solution using Intel® Intelligent Storage Acceleration Library (Intel® ISA-L). Intel® VTune™ Amplifier XE further helped Tencent remove bottlenecks in the TFS system.
- **Intel® Ethernet 10 Gigabit Converged Network Adapter.** As the TFS storage performance improved, the network throughput became a bottleneck for system performance. Tencent migrated to a 10 Gigabit network using Intel® Ethernet 10 Gigabit Converged Network Adapter to resolve the bottleneck.

IMPACT

- **Improving performance, reducing storage costs, and contributing to a green data center.** The erase code solution effectively reduced storage space by 60 percent. After optimization, The I/O performance of TFS improved by 2.8 times and storage performance was enhanced by about 20x. With cold data processed by this system, Tencent has saved significant server resources and raised the performance-price ratio for storage. Meanwhile, this has conserved hundreds of kilowatts of energy for Tencent Data Center, contributing to energy savings and environmental protection.

Tencent is a leader in the Internet industry, with products that have become an indispensable part of people's Internet lives. All these products are supported by Tencent's self-developed file system, TFS.

“TFS is a file service supporting Tencent's businesses,” explained Li Bo, head of the File Storage Team in Tencent's Architecture Platform Department. “Thus, it serves a variety of file types and large volume of data. In fact, Tencent's data volume still experiencing explosive growth. If Tencent tries to solve the problem by simply increasing storage capacity and calculation

nodes, it will face great financial pressure and huge management costs. Obviously, a better approach is to optimize capacity and performance.”

TFS was using a traditional triple-redundancy backup solution to protect the data in the system. Triple-redundancy was a popular data backup solution adopted by the industry. However, as its business grows, Tencent generates a massive amount of new user data every year. Most data—such as images, audios, and videos—will become cold data as time goes by. The waste of storage space in a triple-redundant backup



With a solution based on Intel® VTune™ Amplifier XE and Intel® Intelligent Storage Acceleration Library, Tencent optimized TFS, reducing storage costs and energy consumption, greatly improving product performance, and actively improving the user experience

solution is a real problem for cold data. Compounding the issue, the cost pressure of TFS comes mostly from storage media.

To solve this problem, Tencent engineers first chose to implement an erasure-code solution using the Jerasure* open source library. An erasure-code backup solution transforms data to be stored into a set of parity blocks. This is a solution that optimizes storage space by utilizing computation performance of processors. Tencent decided to adopt a 9+3 approach for optimization of cold data storage in its TFS erasure-code solution. In a 9+3 approach, a piece of data is transformed into a set of 12 parity blocks. A minimum of 9 parity blocks is required to reconstruct the data. That means the data can still be recovered with up to 3 parity blocks missing. Compared with the triple-redundant backup solution, the erasure-code solution significantly reduces storage space while ensuring high availability of data.

During validation of the erasure code solution, Tencent engineers came across a new challenge. They found that the open source Jerasure erasure code solution could only achieve a single core throughput of 41 MB/s to 66 MB/s running on existing Intel Xeon processor-based servers. The throughput was far lower than the I/O throughput of hard drives and network cards, failing to meet Tencent's business demands and also wasting available resources.

Moreover, Tencent's TFS provides online storage services for various products and services, and the corresponding high

volume transactions require extremely high I/O performance. The PathCache* module that is responsible for standard interfaces of file I/O in TFS did not provide the required performance. Collaborating with Intel, Tencent applied a systematic approach to optimization by using Intel's software-performance suite Intel® Parallel Studio XE Professional Edition and Intel architecture platform-based libraries for optimization of calculation performance, Intel® Intelligent Storage Acceleration Library (Intel® ISA-L).

Intel ISA-L provides the tools to help accelerate and optimize storage on Intel architecture systems. Intel ISA-L can run on various Intel® processors and provides operation acceleration through Intel® Advanced Encryption Standard-New Instructions and Intel® Advanced Vector Extensions. In addition, Intel VTune Amplifier XE was used to conveniently identify hot functions and system bottlenecks; detect concurrent synchronization, I/O waiting, and memory leaks; and check the microarchitecture of the CPU to see whether there are performance problems.

Test data from Tencent shows that after field deployment, the optimized erasure code solution actually had its performance doubled but was still limited by the bottleneck of 1Gb Ethernet. Removing the bottleneck with Intel 10Gb Network Interface Cards makes single-core throughput 731MB/s-1109MB/s, over 17 times improvement before optimization.

In addition, PathCache, the module responsible for data I/O interface in TFS, was also optimized by using Intel VTune Amplifier XE. The Intel performance

analyzer helps find bottlenecks in code, identify critical synchronization issues among threads, and speed up communications. The optimized PathCache module performs 2.8 times as fast as the original version, with queries per second (QPS) increasing from 6,000 per second to 17,000 per second.

Since deployment, the optimized TFS has processed tens of petabytes of cold data and effectively reduced storage space by 60 percent. It has saved Tencent from having to purchase approximately 5,000 servers and saved about 500 KW of energy consumption.

"Our data center hardware and software are mainly based on Intel architecture. During TFS optimization for the data center, Intel has given us great support with both its hardware products and software technology. Using consistent hardware and software has also helped us avoid many compatibility problems. We are looking forward to further and extensive hardware and software cooperation with Intel," said Li Bo.

Find the solution that's right for your organization. Contact your Intel representative, visit Intel's [Business Success Stories for IT Managers](#), and check out [IT Center](#), Intel's resource for the IT industry.

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 www.intel.com/performance

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

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.

© 2015 Intel Corporation. All rights reserved. Intel, the Intel logo, Intel Inside, Look Inside, the Look Inside logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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

0515/SHA/PMG/XX/PDF

332517-001EN