CASE STUDY

YAHOO! JAPAN TURNS TO QUOBYTE'S UNIFIED STORAGE SYSTEM TO REDUCE COMPLEXITY AND OPERATIONAL COSTS FOR GROWING STORAGE INFRASTRUCTURE

The internet giant overcomes issues of running infrastructure at scale by unifying file, block, and object storage for low-latency and high-throughput workloads within a single system.

Quobyte

Summary

Needs/Challenges: As one of the world's largest internet companies, Yahoo! JAPAN needed to find the best way to manage growing storage needs while limiting the acquisition and operational expenses that scaling resources require.

Solution: Quobyte's Data Center File System

Platform:

- Quobyte Data Center File System
- OpenStack
- Kubernetes
- 100GbE network
- Commodity servers, SSD, and HDD

Use Case: Internet Service Provider

Key Benefits

- Reduced storage complexity through a unified storage system
- Lower operational expenses, even at massive scale
- Fault tolerance that scales as storage system scales
- Line speed performance for mixed workloads
- Seamless integration with OpenStack
- Native support for Kubernetes
- Easily manage massive storage infrastructure with small team of admins

YAHOO! JAPAN

The Internet Giant Needs to Manage Growing Storage Needs While Reducing Cost and Operational Overhead

Yahoo! JAPAN is one of the world's largest internet companies, currently ranked at No. 28 in Alexa's Top 500 Sites on the Web, with 68 billion pageviews per month. Used by more than 80 percent of Japanese internet users, Yahoo! JAPAN routinely provides access support for more than 90 million devices a day. However, as it sought to evolve from a "smartphone company" to a "data company," Yahoo! JAPAN was challenged to expand its capabilities to leverage vast data stores, including demographic, psychographic, e-commerce, real-time search, and web browsing.

The company operates more than 75,000 physical servers and over 120,000 virtual machines across its six data centers in Japan and one in the United States, with more than 60 PB of storage system capacity total. Many of its services run on Yahoo! JAPAN's private cloud system, which are built on more than 70 OpenStack clusters. These clusters are operated by a team of fewer than 20 engineers.

__The problem__ – Finding the best way to manage growing storage needs while limiting the acquisition and operational expenses that scaling resources require.

In managing clusters of this scale, Yahoo! JAPAN is always looking at ways to reduce operational costs. Key to achieving this goal is implementing solutions that improve efficiency, such as developing a chatbot to help automate support processes. Yahoo! JAPAN has also centralized its system logs and metrics to allow staff to monitor and visualize the system through a single dashboard.

Software-Defined Storage – A Possible Solution

With many private cloud clusters and multiple storage systems dedicated to each cluster, Yahoo! JAPAN was further challenged to reduce overall storage costs. While it uses storage appliances for high availability and to support heavy workloads, they are expensive and not very flexible. The company decided to see if a software-defined storage solution deployed on commodity hardware would provide the flexibility and scalability it needed.

Yahoo! JAPAN implemented its first software-defined storage solution years ago but the challenges of introducing and operating SDS proved more difficult than its deployed storage appliances so it was not fully introduced into production.

"We understood that the network is the most important piece for SDS through past experiences operating SDS products," said Yasuke Sato, infrastructure engineer at Yahoo! JAPAN. "We think the difficult point of SDS is because SDS is a type of distributed system. For appliance storage, internet traffic like rebuilds and generic traffic, control packets for each node goes through high-speed interconnects or backplanes. But for distributed systems, most of the traffic is more of 'east-west' traffic, so we need to have a high-bandwidth, high-availability and low-latency network."

Software Storage for Massively Scalable, Fault-Tolerant Infrastructure

Looking for a solution that could satisfy the need for a massively scalable and fault-tolerant storage infrastructure as it increased its focus on application development and operation, Yahoo! JAPAN came across Quobyte and its next-generation file system. The company began benchmark testing to ensure that the solution would satisfy the performance needs while reducing operational costs, minimizing complexity and providing improved ease of operations.

***The solution** – Quobyte software storage provides support of Yahoo!JAPAN's private cloud infrastructure while providing benefits of massive scalability and fault tolerance.*

Support for Modern Infrastructure – OpenStack and Kubernetes

Backed by a decade of research, Quobyte's Data Center File System serves low-latency and high-throughput workloads within a single system even as the environment grows. By unifying file, block and object storage, Quobyte's parallel file system provides a back end for Yahoo! JAPAN's private cloud, built with OpenStack and Kubernetes containerized infrastructure. A single Quobyte cluster serves as the backend for multiple OpenStack components, including Cinder and Manila, to significantly lower operational costs. Quobyte's proxy service allows high-rate access to file systems through the S3 protocol. Benchmark testing of sequential and random read/write patterns of several block sizes with 18 storage nodes and 70 compute nodes showed that Quobyte provides superior performance to Yahoo! JAPAN's legacy storage system.

"With the Quobyte Data Center File System we can use the all-flash hardware of our choice and are able to achieve the maximum performance of the drives and networking," said Sato. "And the scalable fault tolerance gives us the confidence we need to run as a scalable infrastructure."

Sato said that like with other SDS solutions, there were some issues initially integrating Quobyte into their environment during the evaluation phase, but they were resolved quickly and carefully by Quobyte's support team. This attention to detail provided Yahoo! JAPAN with confidence that it wouldn't have to worry about unresolved problems when implementing Quobyte into its product chain system and proved to be one of the reasons the company chose Quobyte.

The Impact – Storage Infrastructure That Will Easily Scale With the Business

With Quobyte successfully deployed within Yahoo! JAPAN's OpenStack environment, the company is evaluating how best to utilize the unified storage system as part of its future network architecture. Currently, Yahoo! JAPAN's storage nodes have redundant passes using MLAG but it is interested in simplifying with a new configuration where each storage node has a BGP client or by making a client configuration using 100GbE without redundancy, which will both make cabling simpler as well as protecting nodes in the case of a single disk failure. Because of its strong confidence with Quobyte in its product chain system, Yahoo! JAPAN is sure that implementing the SDS solution as part of a future 100GbE or server-side L3 on Clos network configuration will yield additional benefits.

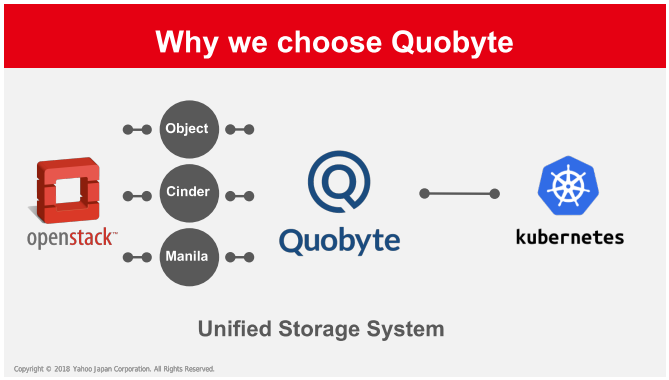

google.com/+Quobyte linkedin.com/company/quobyte twitter.com/quobyte

Quobyte®

Data Center File System.™ Fast and Reliable Software Storage.

Quobyte Inc.
4633 Old Ironsides Drive
Suite 150
Santa Clara, CA 95054
info@quobyte.com
+1 650-564-3111

Quobyte GmbH.
Hardenbergplatz 2
10623 Berlin
info@quobyte.com
+49 30-814 591 800