CUSTOMER SUCCESS STORY

Yahoo! JAPAN adopts NetApp StorageGRID for its platform service, delivering millions of submissions of content daily

NetApp

Internet company updates storage to optimize content delivery platform Castle for container environment and HTTP access

Internet giant Yahoo Japan Corporation wants to leverage cutting-edge technology to make people's lives easier and richer. Yahoo! JAPAN is Japan's largest portal site, offering more than a hundred services, including shopping, auctions, news, videos, weather and disaster information, and maps. More than ten thousand store owners and content partners deliver millions of items of data every day via the Castle content delivery platform. The company is containerizing applications to support these services and decided to review the Castle platform to optimize it for a container environment and HTTP access. The latest-generation Castle platform was launched in October 2020, using the NetApp® StorageGRID® all-flash high-performance object-based storage solution.

YAHOO! JAPAN

Synchronizes data in real time between data centers more than 500km apart.

"As well as breaking down the distance constraints, we can operate StorageGRID as a single storage covering our two data centers in East and West Japan."

Takumi Mizoe
Infrastructure Technology Department 3 Storage 1, Site Operations Division, Yahoo Japan Corporation

Using data analysis to refine Yahoo! JAPAN's services

Starting with the country's largest portal site, Yahoo! JAPAN"s services are evolving at astonishing speed and continuing to create new value. The internet can be leveraged to make life more convenient even for off-line activities like shopping in brick-and-mortar stores. With its PayPay smartphone payment service, Yahoo! JAPAN brings this new experience to everyday life. Koichi Miyata, of the company's Infrastructure Technology Department 2, explains. "Yahoo! JAPAN's Cho PayPay Matsuri event in November 2020 achieved the highest number of users ever. I think this event made people realize how convenient PayPay can be, both online and offline."

Launched by Yahoo! JAPAN and SoftBank in October 2018, PayPay has attracted 33 million users and 2.6 million merchants in just 2 years, rapidly growing into one of Japan's leading cashless payment services. "As well as providing a pleasant cashless payment experience to customers, we want to refine the service by analyzing the huge amounts of data we obtain from it. This will enable us to provide an even better service to customers," says Miyata.

NetApp StorageGRID: High-performance object storage

With Amazon S3 compatible object storage, NetApp StorageGRID offers the scalability to distribute up to 560PB of data to up to 16 data centers as a single namespace. "At Yahoo! JAPAN, we're containerizing applications and moving toward microservices. Step by step, we're optimizing the infrastructure side for container applications too. There's growing demand for user-friendly HTTP-accessible object storage," says Takumi Mizoe. Castle handles all types and sizes of data and files, including images and videos, from several KB to several GB. As soon as Castle receives the data, it is reflected in Yahoo! JAPAN's services: product and stock information for shopping, as well as updates to news and video services.

"Castle is the only data delivery platform between Yahoo! JAPAN and its store owners and content providers. It handles things like weather forecasts, disaster information, and traffic reports that are vital to people's safety and security, so service outages and data loss are not acceptable. When building the new-generation Castle, we set a policy to enhance the structures for service continuity and data protection. One of the specific requirements was that data must be synchronized in real time between data centers more than 500km apart, ensuring service continuity even when a problem occurs at one data center," explains Miyata.

Koichi Miyata
Site Operations Division, Infrastructure Technology Department 2, DC Facility, Yahoo Japan Corporation

Yahoo! JAPAN set the following storage requirements for the new-generation Castle:

1. Amazon S3 compatible object storage allowing HTTP access.
2. Full synchronization in real time between data centers more than 500km apart, to ensure service continuity.
3. Ability to notify systems immediately when data is delivered.
4. Multitenant management for multiple services via API.
5. Capacity and performance easily expandable as processing demands increase.
6. Data that is no longer needed is automatically deleted after a certain time, to maximize use of space.

According to Mizoe, “NetApp StorageGRID met all of our requirements. Being able to introduce the high-performance all-flash appliance (SGF6024) was a major plus point. We also highly rated the ability to configure a huge namespace between two data centers without any distance constraints and being able to flexibly apply data protection policies in a multitenant environment.”

Yahoo! JAPAN prepared hundreds of verification items and spent about 6 months on painstaking testing to evaluate the performance and functions of StorageGRID. They confirmed its compliance with the Amazon S3 protocol and its ability to reliably synchronize between data centers. “We confirmed that even if a NetApp StorageGRID cable or component is removed, it can reliably switch systems to ensure service continuity. It offers outstanding system availability and fault tolerance,” says Miyata.

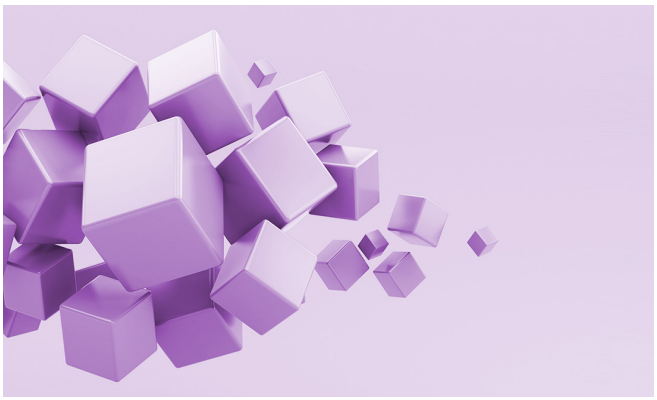

All-flash array allows the system to be used as a data hub

Object storage is generally used to manage large amounts of data at low cost, but Yahoo had a different aim. “Leveraging the high performance of all-flash storage, StorageGRID SGF6024 processes huge amounts of data quickly and reliably. Although it’s object storage, it can handle transactions that require a real-time response,” says Mizoe.

As well as optimal placement of objects between two data centers with ILM, this solution provides a good balance between performance and capacity. “In the existing storage environment there are distance constraints for real-time data synchronization, making it difficult to set up a BCP/DR system for the Tokyo area and Kansai area. As well as breaking down the distance constraints, we can operate StorageGRID as a single storage covering our two data centers in East and West Japan,” says Mizoe.

Takumi Mizoe
Site Operations Division, Infrastructure Technology Department 3, Storage 1,Yahoo Japan Corporation

Flexible application of data protection policies for each service

“We focused on the ability to apply data protection policies using objects and metadata,” says Miyata. The unique ILM policy function of NetApp StorageGRID can optimize the placement, data protection, and duration of objects based on conditions set for each tenant or bucket. According to Miyata, “Using an ILM policy allows us to set the lifecycle, so objects are automatically deleted after a certain period of time. Data placement can also be specified in detail between nodes or between data centers, to ensure service continuity in the event of a node failure or large-scale natural disaster. We should be able to apply ILM policies even more widely.”

Mizoe and Miyata created their own operation monitoring environment for the new-generation Castle. “In addition to the standard monitoring items, we found a way to identify trends in performance and access from a medium- to long-term perspective. NetApp StorageGRID is equipped with Prometheus, which collects the metrics we need, so we were able to implement a monitoring system using Grafana in a short period of time,” says Miyata.

Need for a high-performance storage environment with HTTP access

The creativity of NetApp StorageGRID as a high-performance object storage appliance using ILM policy brought Castle to a new stage. “As we move forward with the process of containerizing applications, there will be an increasing need for a high-performance storage environment with HTTP access. We expect StorageGRID to be applied in a wide range of applications,” says Mizoe. Miyata continues, “Eliminating the distance constraints for real-time data synchronization and being able to cover the data centers that Yahoo uses all over Japan is a huge result from a business point of view. Even if we decide to change data centers, with StorageGRID we won’t need to worry about moving data.”

Mizoe and Miyata conclude, “During the process of implementing optimal data placement using ILM policies, we were able to resolve any issues with the help of NetApp’s technical support. Our communication involving the development team in the U.S. was very smooth, and the physical distance wasn’t an issue. We are truly grateful to NetApp for working on this project with us. We hope they will continue to support our challenges with the latest technology.”

NetApp products
StorageGRID

Protocol
Simple Storage Service

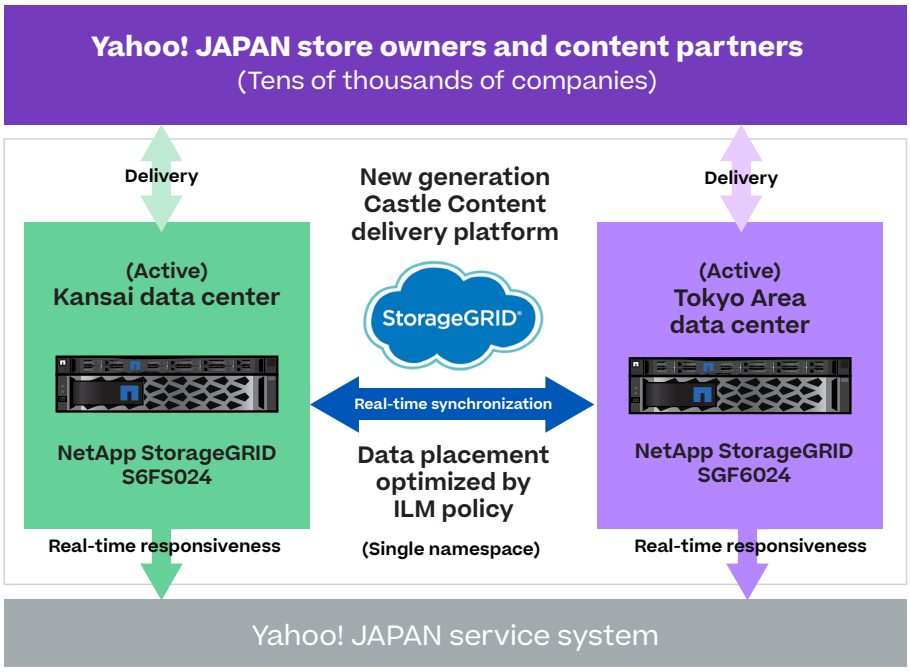

+1 877 263 8277

About NetApp
In a world full of generalists, NetApp is a specialist. We’re focused on one thing, helping your business get the most out of your data. NetApp brings the enterprise-grade data services you rely on into the cloud, and the simple flexibility of cloud into the data center. Our industry-leading solutions work across diverse customer environments and the world’s biggest public clouds.

As a cloud-led, data-centric software company, only NetApp can help build your unique data fabric, simplify and connect your cloud, and securely deliver the right data, services and applications to the right people—anytime, anywhere. To learn more, visit www.netapp.com