

## Yahoo! JAPAN adopted NetApp AFF, FAS, and Astra Trident for stateful services in the world's largest Kubernetes infrastructure



### **Yahoo! JAPAN is promoting stateful applications and the adoption of persistent storage into the Kubernetes-as-a-service infrastructure.**

Yahoo! JAPAN is Japan's largest portal site. The site offers more than a hundred services, including shopping, auctions, news, videos, weather information, disaster information, and maps. The Kubernetes-as-a-service (KaaS) platform that supports this portal site is running more than 860 Kubernetes instances and more than 200,000 containers. In 2018, Z Lab, a wholly owned subsidiary of Yahoo, began improving the environment to provide stateful services on the KaaS platform. Z Lab adopted NetApp® AFF and FAS storage systems, which are highly compatible with Kubernetes and the Container Storage Interface (CSI), and the reliability of the compatibility has been evaluated. In the KaaS platform that supports a secure multitenant environment, NetApp's CSI-compliant storage orchestrator, NetApp Astra™ Trident, enables dynamic storage provisioning and improved storage controls.

## (KaaS) platform runs more than 860 Kubernetes instances

“We chose NetApp AFF and FAS as the optimal storage for our KaaS platform and adopted the CSI-compliant Trident tool provided by NetApp.”

Mr. Aki Numata  
Leader, Storage Department 1, Site Operations Division,  
Technology Group, Yahoo Japan Corporation


### Improved execution environments for stateful applications

The private cloud platform that supports the various services of Yahoo! JAPAN has undergone major changes by incorporating cloud-native technology. Yahoo! JAPAN's stateful applications in Kubernetes as a service (KaaS) have been available for service departments across the company since 2018. Z Lab is responsible for the research and development of next-generation infrastructure.

“The private cloud platform that supports the various services of Yahoo! JAPAN have already been built into containers and microservices,” said Yukinori Sakashita of Z Lab. “As of the end of 2020, KaaS, which is a unique managed Kubernetes service, has more than 860 Kubernetes instances and more than 200,000 containers in operation, and is used in more than 200 services. Being able to manage storage on Kubernetes with stateful services is of great significance to users.”

A bird's eye view of the KaaS platform of Yahoo! JAPAN shows that Kubernetes is used to control not only the container but also the virtual machine on which Kubernetes is set up.

“By using Kubernetes to manage Kubernetes, the strengths of Kubernetes such as self-healing, rolling updates, and scaling are applied to the entire KaaS platform,” said Mr. Sakashita. “This promotes the autonomous operation of more than 860 Kubernetes instances, enabling the operation of a large-scale KaaS platform with more than 200,000 containers operating. It is also notable that the virtual machine on which Kubernetes is set up is modeled so that it can be treated as the same resource as the container, and the operation of many Kubernetes instances is labor saving according to the framework of Kubernetes.”



“At Yahoo! JAPAN, stateless and operable applications such as web applications had been containerized without having data managed by Kubernetes. On the other hand, stateful applications that handle data had been provided as a hybrid type that is managed by tools other than Kubernetes.”

Aki Numata, Site Operations Division, System Management Division, Yahoo Technology Group, said, “In response to an internal request to use stateful applications in a KaaS environment, we started to consider persistent storage on the KaaS platform in 2018. At the time of consideration, we emphasized high reliability that does not stop the service of Yahoo! JAPAN, which plays a role as a social multitenant infrastructure system that securely integrates multiple services while conforming to Kubernetes and bringing speed to the business.”

### Building persistent storage with NetApp Astra Trident

Yahoo! JAPAN has a policy of prioritizing storage according to the importance of data. For this reason, they chose NetApp storage, which has a proven track record and is highly reliable for data of high importance.

“We thought that a proven storage appliance would be the best way to achieve a highly reliable multitenant environment. We chose NetApp AFF and FAS as the optimal storage for our KaaS platform and adopted the CSI-compliant Trident tool provided by NetApp,” said Mr. Numata.

The NetApp AFF A700 all-flash array and NetApp FAS9000 hybrid array adopted for persistent storage provide the optimum storage environment for the KaaS platform of Yahoo! JAPAN, leveraging the various functions of the NetApp ONTAP® data management software. In addition, ONTAP provides flexible control of NetApp storage from Kubernetes via Astra Trident.

Mr. Sakashita emphasized the importance of speed. “In 2015, an infrastructure renewal project was launched with the aim of delivering the service value of Yahoo! JAPAN to customers as soon as possible. The theme raised here is speed,” he said.

“To that end, we have developed containerization and microservices, developed CI /CD development methods, and promoted support for cloud-native technology. Under these circumstances, Z Lab is working on container technology to support cloud-native tools and is developing and supporting the KaaS platform.”

Yahoo! JAPAN’s KaaS infrastructure has the capacity to support faster development. The provisioning of persistent storage with NetApp storage and Astra Trident will bring similar speed to stateful applications.

“The key to speed is to make storage management cloud native by controlling NetApp storage from Kubernetes via CSI drivers,” said Mr. Numata. “It’s important to provide an environment where service developers can easily allocate and expand volumes when they need them.”

### **Astra Trident evolving at the same speed as Kubernetes**

Astra Trident is provided by NetApp as open-source software. It is a CSI driver for using NetApp storage from Kubernetes, and it has a full range of functions as a storage orchestrator and provisioner.

“Kubernetes continues to evolve at an alarming rate at the heart of the cloud-native ecosystem,” said Mr. Sakashita. “Trident’s strength is that it is updated at the same speed as Kubernetes. Following Kubernetes, which releases the latest version every 3 months, we are steadily enhancing the functions required for cloud-native operation. Our request for the addition of functions was also met at a high speed.”

One example of a request presented by Yahoo/Z Lab is the development of a function to automatically update the Access Control List (ACL) on the storage side, along with rolling updates and self-healing.

“Automatic storage ACL updates are not covered by the CSI standard,” said Mr. Sakashita. “However, if [it relies on] the manual work of the storage administrator, it will not be a cloud-native operation. A few months after the request, NetApp’s feature updates added automatic ACL updates.”

In the process of realizing a stateful application environment, some issues became apparent, but Yahoo/Z Lab worked with NetApp to resolve them and improve the capabilities of Astra Trident as a CSI driver.

“Like the open-source community, it was great that we were able to form a team with the idea of “Let’s brush up together,” said Mr. Numata. “The support of the NetApp Japan team was also timely and appropriate.”

### **Huge multitenant environment to operate various services**

In the summer of 2020, Yahoo! JAPAN began offering a KaaS platform that supports stateful applications. Users select Persistent Storage by NetApp AFF and FAS from their own dedicated portal developed by Yahoo and create a tenant storage virtual machine (SVM) and storage pool through the approval process. This provides Kubernetes with independent storage for each user — for example, a service department.

According to Mr. Numata, “For persistent storage with NetApp, the SVM provided by the ONTAP data management software allows us to configure completely independent tenants on a per-service-department basis. By preparing a virtual storage controller for each service department, the network is separated. It is a secure environment where access is limited to the storage resources prepared for each service. Because we operate a system that plays a part in social infrastructure, we aim for 100% data availability for storage systems. Yahoo uses NetApp’s data protection features such as SyncMirror and SnapMirror to protect data inside and outside the site according to our requirements. With NetApp, we get the data availability and cloud-native operation required for social infrastructure systems. In order to achieve both, we cross technologies and organizations as necessary, regardless of existing mechanisms and systems. It is necessary to further expand, and it is important to select the optimum storage according to the characteristics of the application and data.”


“We chose NetApp storage because of its track record and reliability, and it provides a stable persistent storage environment as expected,” said Mr. Sakashita. “Trident, on the other hand, is under aggressive development. Stability and aggression—not many storage vendors can meet these two demands at the same time. I hope NetApp will continue to take on the challenge of becoming a cloud-native provider without slowing down its current sense of speed.”

## NetApp products

- Astra Trident
- NetApp AFF
- NetApp FAS
- ONTAP
- SnapMirror
- SyncMirror

## Protocols

- NFS
- iSCSI



+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](http://www.netapp.com)



© 2022 NetApp, Inc. All Rights Reserved. NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners. CSS-7190-0222