

Rakuten Cloud Builds Stateful Edge Cloud Networks with Intel

Rakuten Cloud-Native Platform and Intel® architecture-based edge technologies provide low latency, reduced costs and scalability



One of the challenges of deploying edge cloud networks is supporting a growing number of stateful applications using Kubernetes, which is designed for stateless workloads.

An application’s state is its status at any point in time. A stateful application saves this state information to storage and can act or build upon it. Other applications are considered stateless in that once an action or user request has taken place the application no longer keeps track of the state of that transaction. Each follow on action or transaction is executed without regard for the previous state information.

Rakuten Symphony

Many web applications are stateless but offer some stateful capabilities by offloading these requests to a user’s device or to a database server. For example, a web browser is a stateless application, but with cookies stored on the user’s device, or with access to a database, a browser can deliver stateful features.

Enterprises are expanding their use of edge cloud servers that need to support stateful applications such as databases, analytics, IoT, mobile networks, AI model training and others. Many edge use cases, such as computer vision or analytics, can benefit from low latency access to local data at the edge which makes stateful edge application delivery a necessity. The volume of data processed and the need for low latency limits the ability to offload the data to an out-of-container storage solution. This makes support for stateful cloud-native applications essential.

Kubernetes is the leading cloud-native operating system and was designed to orchestrate stateless applications. The main challenge to building stateful cloud-native applications is delivering persistent data storage within a container or a Kubernetes pod to maintain state information.

Many initial cloud deployments relied on managed database services or ran their databases outside of Kubernetes containers / pods because round-trip times impacted latency, reducing performance and requiring new reliability and reliance infrastructure for the storage.

Table of Contents

- Other Edge Cloud Challenges 1
- Rakuten Cloud Provides Complete Stateful Edge Solution.....2
- Rakuten Cloud-Native Platform Optimized for the Edge2
- Simplified Workload Placement.. 2
- Services Orchestration and Lifecycle Management.....2
- Works with All Deployment Modes.....3
- Cloud-Native Storage for Stateful Applications3
- Application-Aware Storage3
- Servers based on 4th Gen Intel® Xeon® Scalable Processors.....3
- Conclusion.....4
- Learn More.....4

Other Edge Cloud Challenges

In addition to supporting more stateful applications, edge cloud deployments are challenged by the remote placement of the servers. Edge cloud servers are located at remote offices and other distant locations that are close to the customer but away from the network manager. This makes automation and orchestration important for cost-effective management of the servers and reduces the need for onsite maintenance requests.

Achieving remote orchestration and management starts with integrating Kubernetes with the server and testing the functionality. Kubernetes is the foundation for edge cloud infrastructure and must be tightly integrated with the servers for general operations and special features such as application placement and resource inventory. Testing the integration is important for problem-free remote management.

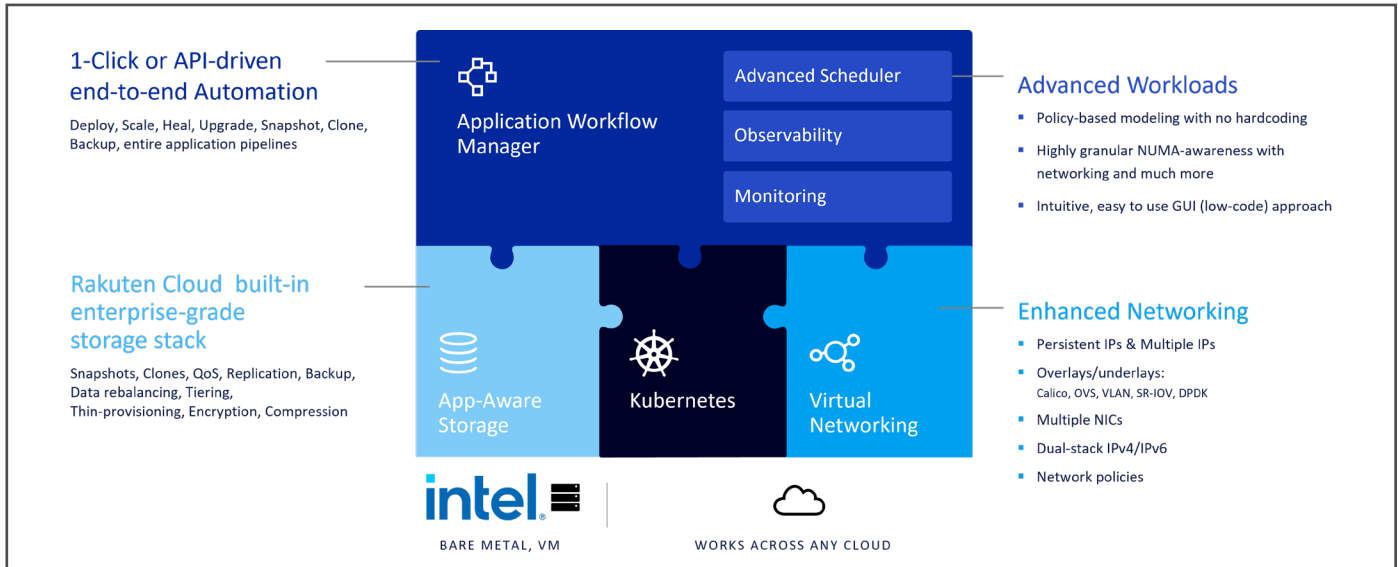


Figure 1. Block diagram shows the three key software components of Rakuten Cloud solution.

Rakuten Cloud is the cloud business unit of Rakuten Symphony, an Intel® Network Builders Titanium level community member. Rakuten Cloud is a technology leader in building stateful edge cloud to help organizations simplify edge cloud deployments and life cycle management. The Rakuten Cloud software suite is tightly integrated with servers running 4th Gen Intel® Xeon® Scalable processors.

Rakuten Cloud Provides Complete Stateful Edge Solution

Rakuten Cloud’s cloud-native environment has three main applications (Figure 1) including the Rakuten Cloud-Native Platform (CNP), Cloud-Native Orchestrator (CNO) and Cloud-Native Storage (CNS).

Rakuten Cloud-Native Platform Optimized for the Edge

Rakuten CNP is an industry-leading Kubernetes platform that is optimized for running storage and network-intensive applications in edge cloud environments. The software was built with a small footprint for edge applications – requiring only two CPU cores for full functionality and observability.

Rakuten CNP is ideal for organizations looking to deploy AI/ML, retail, Industry 4.0, massive IoT, financial services, private 4G/5G and other similar applications.

For flexible resource pooling, CNP supports both container and virtual machine-based network functions for resource pooling across all the same cluster or separate clusters. This support enables orchestration workflows spanning physical network functions (NFs), virtual NFs and container NFs.

Rakuten CNP provides an Application Workflow Manager layer on top of Kubernetes that can onboard new applications using a single-click interface or by using an API. The software can deploy a complete application pipeline seamlessly.

Simplified Workload Placement

Rakuten Cloud-Native Platform features an advanced declarative workload placement capability that allows users to ask for required resources and then an automated resource reservation system finds those resources – such as CPUs, NICs, persistent IP addresses – reserving them for the application.

This declarative interface is ideal for both day one service turn up activities, as well as ongoing planning and lifecycle operations. It also provides a better understanding of failover behavior before it happens, making more efficient use of resources and reducing human error.

The comprehensive workload placement algorithm auto detects, connects, and configures all of the resources needed for the application based on resource policies.

Popular Rakuten CNP placement variables include NUMA-awareness, CPU-pinning, multi-service affinity / anti-affinity policies, min/max IOP values to eliminate noisy neighbors, auto-discover SR-IOV-enabled NIC cards, FPGA and GPU resources from the same NUMA node.

Services Orchestration and Lifecycle Management

The Rakuten Cloud-Native Orchestrator (CNO) is a server and application management automation platform that enables edge cloud scale. The software supports bare-metal servers, third-party applications, cloud native containers, and virtual machines across the entire stack from edge to core. This enables seamless deployment and management of diverse, secure edge computing capabilities.

For automation, Rakuten CNO has built-in data collection capabilities including logging, monitoring and a policy engine that drives the closed-loop automation capability. It enables unified orchestration of containerized, virtualized and bare-

metal workloads across distributed edge clouds using a GitOps model. It also supports multi-tenancy, role-based access control and resource pooling for secure collaboration.

A vendor ecosystem is part of Rakuten CNO allowing MNOs to onboard both enterprise and telecom applications from a large and growing number of third-party vendors with one click. This ensures that MNOs have the most innovative solutions and simplifies the massive job of rolling out big networks spanning multiple locations. The system can also ingest scripting to automate methods of procedures (MOPs) for the orchestration of any network element or process.

Rakuten CNO facilitates advanced inventory and service management across multiple clusters with an automated policy engine that intuitively relates layers of dependencies to know their impact on services. For example, an MNO could drill down into an underperforming cluster or hardware node to see what network elements are degraded or have failed. The impact of those faults on any device, node or process can be analyzed in order to model the impact of a failure.

Works with All Deployment Modes

Rakuten Cloud-Native Orchestrator has lifecycle management capabilities for bare metal servers, cloud platforms / Kubernetes, and applications / network functions. For all of these server deployment modes, Rakuten CNO offers remote software and firmware downloads and configuration.

For bare metal implementations, Rakuten CNO features numerous intent-driven and contextually aware checks to help develop declarations for management. Rakuten CNO also allows remote server transformation including verification, installation, upgrading or patching the OS, and adding and configuring drivers, services and software patterns.

For cloud platform / Kubernetes clusters, Rakuten CNO has zero-touch automation of Kubernetes cluster deployments across thousands of edge cloud locations. Rakuten CNO's close integration with the Kubernetes-based Rakuten Cloud-Native Platform delivers features such as resource modeling and advanced workload placement based on NUMA-awareness, CPU pinning, huge pages, affinity / anti affinity rules and multi container runtime interfaces (CRIs) for both containers and virtual machines.

For applications and network functions, Rakuten CNO is able to manage both VMs and containers on the same cluster for resource pooling. These applications also benefit from the workload placement features as well as networking and application aware storage from Rakuten Cloud-Native Storage (CNS).

Cloud-Native Storage for Stateful Applications

Developing a stateful cloud-native edge network requires adding a storage layer to Kubernetes that is aware of Kubernetes and applications in order to accelerate data management, data protection and application deployment times.

Rakuten CNS is designed to handle the microservices architecture of Kubernetes. Microservice apps have an altered relationship to data because of their segmented approach to app design. CNS alleviates this challenge by working in the context of the entire application by managing lifecycle and data protection operations for the entire application including data, metadata, and Kubernetes configurations.

The software creates a local storage abstraction layer on each edge node, making this local storage available to workloads running on other edge nodes.

This software defined storage (SDS) solution offers advanced application awareness, high availability, data resiliency, and security that enables customers to migrate their stateful workloads to containers seamlessly and transform their edge deployments.

Rakuten Cloud-Native Platform also has a full slate of storage features (see Table 1) including snapshots, clones, replication, backup, data rebalancing, tiering, thin provisioning, encryption, and compression.

Application-Aware Storage

Rakuten CNS can work with Kubernetes to allocate storage during deployment of an application or cluster and share storage among apps and users.

Rakuten CNS enforces key service level agreement (SLA) guarantees when consolidating applications and features, it also supports data locality, affinity, anti-affinity and isolation constraints. The software can facilitate applications that need to modify the root filesystem, create snapshots. It also supports cloning, QoS, replication, backup, data rebalancing, tiering, thin-provisioning, encryption and compression.

Servers based on 4th Gen Intel® Xeon® Scalable Processors

Edge cloud servers need compute capacity and performance to handle advanced applications today and to be future-proof for emerging applications. Rakuten Cloud Cloud-Native Software suite works with servers powered by 4th Gen Intel Xeon Scalable processors.

Any stateful app	RDBMS, NoSQL, Queues, Timeseries, Analytics
Storage & data management	Rakuten Cloud-Native Storage
Any Kubernetes	Open source Kubernetes, Anthos, OpenShift, GKE, AKS, EKS, IKS, RKE, Rakuten CNP
Any infrastructure	Bare metal, VMs, AWS, Azure, Google Cloud, IBM

Table 1. Rakuten Cloud-Native Storage features.

This CPU family offers high-throughput and low-latency with devices that are engineered with up to 60 cores (120 cores with Intel® Hyper-Threading Technology (Intel® HT Technology) for high performance edge cloud deployments. The processor family's architecture combines high-performance processor cores with up to eight built-in accelerators for maximum performance efficiency on key workloads. Integration of accelerators into the processor redefines CPU architecture and provides a more efficient way to achieve higher performance than relying solely on increasing the CPU core count for workload processing.

In 5G radio access network (RAN) use cases, the Rakuten Cloud solution makes use of the Intel® vRAN Boost accelerator that offloads computationally heavy layer 1 tasks such as low-density parity check (LDPC) decoding and forward error correction (FEC).

4th Gen Intel Xeon Scalable processors offer up to 80 lanes of PCIe 5.0 connectivity and support Compute Express Link (CXL), a cache-coherent interconnect for processors, memory expansion, and accelerators. Another important accelerator for the Rakuten Cloud distributed stateful edge is support for virtualized connections to physical PCIe devices using single root IO virtualization (SR-IOV) and/ or Intel® Scalable I/O Virtualization (Intel® Scalable IOV).

Benefits of Stateful Edge Cloud

- Low Latency
- Reduced Bandwidth Costs
- Data Sovereignty and Security
- Scalable Cloud Capabilities at the Edge
- Automation and Simplified Operations

With Intel® Advance Matrix Extensions (Intel® AMX), the 4th Gen Intel Xeon Scalable processors have exceptional AI training and inference performance. Other seamlessly integrated accelerators include Intel® QuickAssist Technology (Intel® QAT) that speeds up data encryption/decryption and compression for faster query throughput for more responsive analytics.

Conclusion

Supporting the emerging stateful edge cloud requires an infrastructure layer that offers flexibility, automation, and stateful storage. Rakuten Cloud-Native Software suite is built on Kubernetes to offer cutting-edge containerization and enhances that orchestration and cloud-native storage.

Combined, this software suite provides a complete edge network solution with the one-click deployment and advanced workload placement features needed for edge cloud deployments that span hundreds or thousands of servers.

For performance, the solution runs on servers powered by 4th Gen Intel Xeon Scalable processors that offer compute performance, large core count and on-chip accelerators to accommodate the wide range of applications that will be run in the edge cloud. Rakuten Cloud and Intel make it easy to build an edge cloud network to support your enterprise.

Learn More

[Rakuten Cloud](#)

[4th Gen Intel® Xeon® Scalable processor](#)

[Intel® Network Builders Ecosystem](#)

[Intel® Industry Solution Builders](#)



¹<https://www.intel.com/content/www/us/en/products/details/processors/xeon/scalable.html>

Notices & Disclaimers

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.