

Rackspace Spot: Kubernetes Control Planes as Workloads, at Scale

**How Rackspace built a turnkey
managed Kubernetes service on Kamaji
– spinning up isolated control planes in
seconds and running hundreds on
shared infrastructure**

When Every Cluster Needs Its Own Control Plane, Scale Gets Expensive

COMPANY SUMMARY

- **Company name:** Rackspace Technology (NASDAQ: RXT)
- **Website:** www.rackspace.com
- **Headquarters:** Texas, USA
- **IT Environment:** Multi-tenant managed Kubernetes across multiple regions and clouds; Hosted Control Planes (Kamaji) on a shared management cluster; dynamic spot/auction compute for worker nodes
- **Industry:** Hybrid cloud and AI solutions / managed cloud services
- **No. of employees:** ~5,000

HIGHLIGHTS


- Fully isolated control planes spun up in seconds, not hours
- Hundreds of control planes hosted on a single shared management cluster
- Worker nodes anywhere: multi-region, multi-cloud, connected via Konnectivity
- One consistent lifecycle (upgrades, patching, scaling) across every cluster
- Global, elastic managed Kubernetes at a fraction of the cost

Rackspace Technology (NASDAQ: RXT) is a leading end-to-end hybrid cloud and AI solutions company. With Spot, Rackspace set out to deliver fully managed Kubernetes clusters at compelling cost-efficiency, powered by dynamic spot/auction compute and delivered as a turnkey experience.

For a platform like Spot, keeping Kubernetes operational is not a background concern – it is the product. Every customer running on Spot expects clusters on demand, predictable performance, and complete lifecycle management including upgrades, patching, and failover, with each tenant experiencing an isolated, single-tenant-quality environment, even while sharing underlying infrastructure. Delivering that reliably, and at the scale of hundreds or thousands of clusters, is the promise Rackspace makes to its customers.

CHALLENGES

To make an auction-based instance marketplace work, Rackspace had to satisfy three goals at once: cost efficiency, rapid provisioning and elasticity, and operational consistency and reliability. Under the traditional Kubernetes model, those goals conflict.



The conventional approach dedicates three VMs or hosts to every cluster's control plane – etcd, API server, and controllers. That adds a fixed cost per cluster that simply does not scale when the target is hundreds of clusters.

Manual or semi-automated provisioning is far too slow for short-lived spot instances, where customers expect clusters to appear in seconds rather than minutes. And once the fleet grows, per-cluster operations – upgrades, patching, monitoring, logging – become an operational nightmare. Rackspace needed a new model.

SOLUTION: Control Planes as workloads with Kamaji

CLASTIX's open-source project Kamaji (Apache 2.0) treats Kubernetes Control Planes as pods – as workloads – running inside a shared management cluster. This is the Hosted Control Plane (HCP) model, and it maps directly onto what Spot needed.

With Kamaji, each tenant cluster's control plane (API server, controller-manager, scheduler, and datastore) runs as a set of pods in the management cluster. Lifecycle operations – creation, scaling, upgrade, and deletion – become as simple as applying a KamajiControlPlane resource through Cluster API.

Because many control planes share the same infrastructure, Rackspace achieves high density without sacrificing isolation: each tenant is separated by namespaces, RBAC, and network segmentation, with flexible datastore options (etcd or MySQL, shared or dedicated) to match performance and isolation requirements. Everything is Kubernetes-native, so Rackspace reuses its existing automation – operators, CRDs, GitOps, and Cluster API.

Spot's worker nodes can live anywhere: across regions, across availability zones, even across clouds. To connect them to a centrally managed control plane, Rackspace relies on Connectivity. Worker nodes initiate outbound connections to the API server, so no inbound access into worker nodes is required. This keeps workers firewalled and isolated, simplifies network topology, and keeps latency low – typically in the tens to low hundreds of milliseconds, negligible for microservices workloads. From the workload's perspective, the control plane behaves as if it were next door, wherever it physically sits.



“We're confident in recommending this architecture as the foundation for the next generation of managed Kubernetes services.”

–Kevin Carter, Director, Open Source Product

Ongoing Benefits: Scalability, Efficiency, Uniformity, Elasticity... and Lower Costs

By moving from "one cluster equals three control plane VMs plus workers" to control planes as workloads, Rackspace unlocked benefits across cost, speed, and operations:

- **Scalability** – control planes spin up in seconds, letting the service grow without a linear increase in operational overhead.
- **Resource efficiency** – a single management cluster hosts hundreds of control planes on shared infrastructure, instead of a dedicated VM set per tenant – making far better use of hardware and underpinning Spot's low-cost model.
- **Operational uniformity** – upgrades, patching, monitoring, and logging follow one consistent pattern across every control plane, replacing per-cluster custom scripts.
- **Elasticity** – control plane pods autoscale on a rolling basis, idle clusters can be spun down, and everything is managed declaratively through APIs and infrastructure-as-code.
- **Global reach at lower cost** – because workers connect outbound via Konnectivity, Rackspace does not replicate control planes per region, delivering a top-tier service at a fraction of the cost.

For Spot's customers, the result is faster provisioning, predictable performance, consistent lifecycle management, and a Kubernetes experience that is genuinely global and elastic. Rackspace has publicly endorsed Kamaji, with Kevin Carter, Director of Open Source Product at Rackspace Technology, describing the architecture as "the foundation for the next generation of managed Kubernetes services."



Source: The information in this case study is drawn from Rackspace's publicly published blog post, "A new paradigm for cloud-native infrastructure" (24 November 2025), available at: https://blog.rackspacecloud.com/blog/2025/11/24/a_new_paradigm_for_cloud-native_infrastructure

Employee count estimate is per Rackspace's FY2025 10-K, as of 31 Dec 2025

NEXT STEPS

If your team is building a managed Kubernetes service, a multi-tenant platform, or a sovereign or edge offering, Kamaji lets you run control planes as workloads – dense, elastic, and centrally managed, wherever your worker nodes live. Reach out to the CLASTIX team, and we will help you find the right architecture for your environment, without starting from zero.

About Rackspace Technology

Rackspace Technology® (NASDAQ: RXT) is the operator of the full enterprise AI stack from governed private cloud to AI inference and agents in production. With an Outcomes-as-a-Service model built on secure infrastructure, data foundations, and forward-deployed engineering, Rackspace delivers business results for regulated and mission-critical industries where governance, sovereignty, and uptime are non-negotiable. Learn more at www.rackspace.com.

About CLASTIX

Clastix builds enterprise-grade Kubernetes infrastructure through hosted control planes and hard multi-tenancy – solving cluster sprawl, fragile isolation, and runaway infrastructure costs at scale. With deep open source credibility – including the donation of Capsule to the CNCF – the kMetal platform is purpose-built for cloud providers, MSPs, and platform teams that demand real isolation, efficiency, and scale. Customers include NVIDIA and Mistral AI.

Contact us

Please contact us to request a quote or learn more about our support offerings.



CLASTIX SRL
V. Francesco Caracciolo, 17
80122 Naples, Italy

VAT No. IT11183320966
REA: NA-1152378
© 2026 Clastix | All Rights Reserved

hello@clastix.io
www.clastix.io
<https://github.com/clastix>