

# Server Leftovers

By *Roy Schestowitz*

Created *13/11/2019 - 11:36pm*

Submitted by Roy Schestowitz on Wednesday 13th of November 2019 11:36:15 PM Filed under [Server](#) [1]

- [Knative at 1: New Changes, New Opportunities](#) [2]

This summer marked the one-year anniversary of Knative, an open-source project that provides the fundamental building blocks for serverless workloads in Kubernetes. In its relatively short life (so far), Knative is already delivering on its promise to boost organizations' ability to leverage serverless and FaaS (functions as a service).

Knative isn't the only serverless offering for Kubernetes, but it has become a de-facto standard because it arguably has a richer set of features and can be integrated more smoothly than the competition. And the Knative project continues to evolve to address businesses' changing needs. In the last year alone, the platform has seen many improvements, giving organizations looking to expand their use of Kubernetes through serverless new choices, new considerations and new opportunities.

- [Redis Labs Leverages Kubernetes to Automate Database Recovery](#) [3]

Redis Labs today announced it has enhanced the Operator software for deploying its database on Kubernetes clusters to include an automatic cluster recovery that enables customers to manage a stateful service as if it were stateless.

Announced at Redis Day, the latest version of Kubernetes Operator for Redis Enterprise makes it possible to spin up a new instance of a Redis database in minutes.

Howard Ting, chief marketing officer for Redis Labs, says as Kubernetes has continued to gain traction, it became apparent that IT organizations need tools to provision Redis Enterprise for Kubernetes clusters. That requirement led Redis Labs to embrace Operator software for Kubernetes developed by CoreOS, which has since been acquired by Red Hat. IT teams can

either opt to recover databases manually using Kubernetes Operator or configure the tool to recover databases automatically anytime a database goes offline. In either case, he says, all datasets are loaded and balanced across the cluster without any need for manual workflows.

- [Dare to Transform IT with SUSE Global Services](#) [4]

[Server](#)

---

**Source URL:** <http://www.tuxmachines.org/node/130470>

**Links:**

[1] <http://www.tuxmachines.org/taxonomy/term/147>

[2] <https://www.tfir.io/knative-at-1-new-changes-new-opportunities/>

[3] <https://containerjournal.com/topics/container-ecosystems/redis-labs-leverages-kubernetes-to-automate-database-recovery/>

[4] <https://www.suse.com/c/transformation-with-suse-global-services/>