

Automate pod rightsizing at scale

Challenge

While VPA adoption is low due to pod restarts, imprecise scaling, and HPA incompatibility, manual vertical scaling requires continuous monitoring, and can cause CPU throttling, downtime, OOM errors, or inflated costs in case of overprovisioning.

Solution

Zesty's Pod Rightsizing automatically adjusts CPU & memory resources at the container level to match precise, real-time workload usage, maximizing stability & performance while reducing costs.

Benefits

- Ensure performance & stability
- Reduce costs
- Eliminate manual monitoring
- Container-level VPA
- Compatible with HPA

How it works

- 1 Zesty is granted IAM role access permission
- 2 Zesty installs a Kubernetes agent with a read-only-permission
- 3 You connect your CUR to Zesty platform & install Pod Rightsizing
- 4 Zesty collects and analyzes workload utilization and pattern histories
- 5 Zesty Pod Rightsizing generates recommendations to optimize resource allocation and generate savings
- 6 Once you activate the recommendation, Zesty scales up or down the container's request

