



Headroom reduction

zesty

Kubernetes Optimization
Platform

Scale nodes faster to enhance efficiency

Challenge

Deploying a new node in a Kubernetes environment can take several minutes. To overcome this limitation and ensure application availability during traffic peaks, DevOps teams often over-provision nodes, leading to inefficiencies and unnecessary costs.

Solution

Zesty's HiberScale™ technology accelerates Karpenter and Cluster Autoscaler scaling to under 30 seconds, eliminating the need for CPU overprovisioning and reducing costs, while maintaining SLAs during unexpected traffic spikes.

Benefits



Reduce node headroom



Cut cluster costs by 70%




Eliminate manual operations



Preserve SLAs

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 the Zesty agent, which starts collecting workload usage and pattern histories
- 4 Zesty analyzes your system's scaling performance, identifies unnecessary headroom
- 5 You install the K8s scaler, which automatically creates a pool of hibernated nodes, minimizing headroom
- 6 In your Kubernetes environment, you reduce your replicas to the minimum recommended level
- 7 The hibernated nodes are re-activated within 30 seconds when needed, in response to spikes or increased CPU requests



zesty supports  

Contact Us!  www.zesty.co  info@zesty.co

