Headroom reduction

zesty

Eliminate Node Headroom With Zero SLA Compromise

Challenge

Deploying a new node in a Kubernetes environment can take 5 minutes and sometimes even more. 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 unique HiberScale technology enables node deployment 5X faster, eliminating the need for a node buffer during traffic peaks.
This automatically reduces costs without compromising SLAs.

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

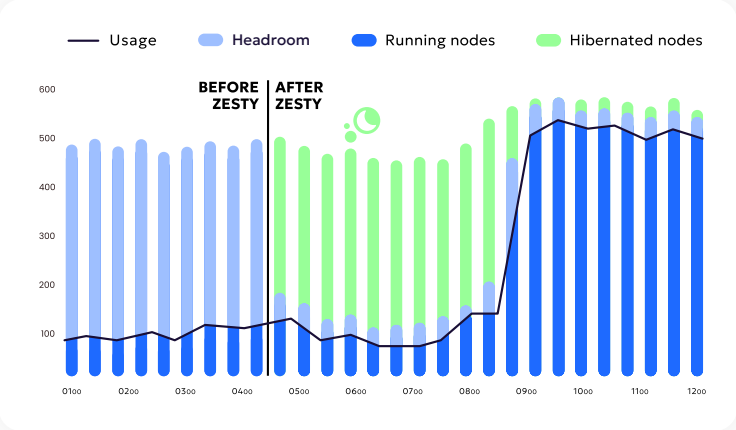

4. Zesty analyzes your system's scaling performance, identifies unnecessary headroom
5. In your Kubernetes environment, you reduce your replicas to the minimum recommended level
6. You install the K8s scaler, which automatically creates a pool of hibernated nodes, minimizing headroom
7. The hibernated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