Kompass

Zesty Automated K8s Optimization Platform

zesty

Cut Kubernetes costs with zero SLA compromise.

Leverage multi-dimensional automation with a suite of solutions designed for effortless, continuous and durable Kubernetes cost optimization.

The problem

Optimizing costs while ensuring 100% availability

Managing and optimizing a dynamic Kubernetes environment is a complex task. DevOps teams' main challenge is to ensure application resiliency and maintain SLAs, while optimizing costs.

Limitations in compute and storage resource scaling push teams to over-provision, leading to unnecessary expenses and missed savings opportunities.

Zesty's solution

Automated Kubernetes optimization platform

Zesty Kompass provides a holistic suite of solutions that streamline cost reduction while preserving resiliency and SLAs.

With a multidimensional automation approach, and unique technologies designed for nodes and persistent volumes auto-scaling, Zesty's automated Kubernetes optimization platform ensures effortless and continuous optimization of Kubernetes resources with real-time needs.

Benefits

Save up to 70%
Cut Kubernetes costs through automated resource optimization.

Preserve SLAs & resiliency
Ensure application resiliency and maintain SLAs with real-time autoscaling.

Unify Kubernetes optimization
Manage all Kubernetes optimizations in one single platform.

Boost operational efficiency
Leverage automation to reduce manual intervention and optimize workloads.

Kompass

Zesty Automated K8s Optimization Platform

zesty

What makes our platform different?

Multi-dimensional automation

Zesty developed a unique multi-dimensional approach, combining multiple automated solutions across both compute and storage: tailored recommendations, node headroom reduction, Spot automation, and Persistent Volumes autoscaling.

This comprehensive approach ensures continuous, effortless, and durable Kubernetes cost optimization, delivering unmatched savings, with no compromise on resiliency and SLAs.

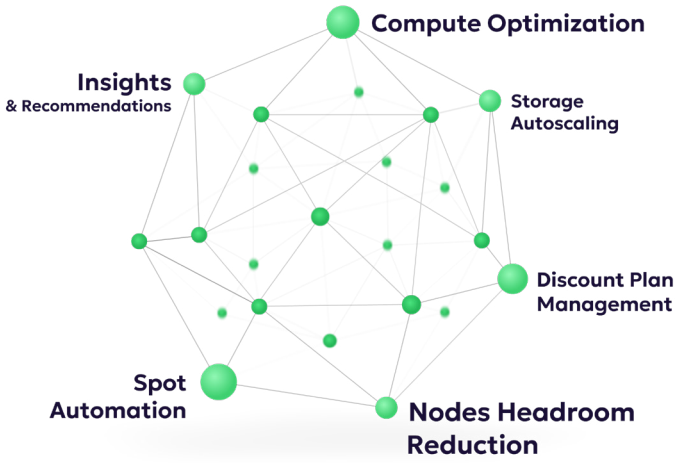

HiberScale technology

Zesty's HiberScale technology enables large-scale node hibernation, and their instant reactivation within 30 seconds, in case of spike. This unique capability enhances cost efficiency without compromising SLAs.

Without Zesty, a considerable node headroom is created to maintain SLAs, leading to **unnecessary costs**.

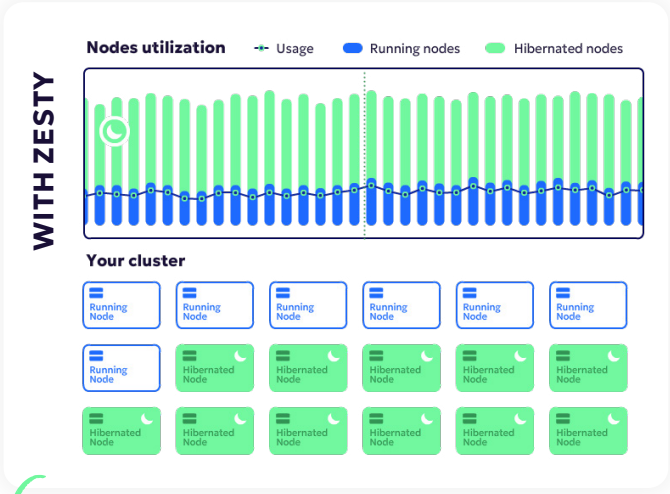

Zesty creates a pool of **hibernated nodes,** reducing headroom to minimum.

Insights

zesty

Identify Kubernetes savings opportunities

Challenge

Having clear visibility into costs and resource utilization is a persistent challenge in dynamic, complex Kubernetes environments.

Solution

Insight provides granular visibility over clusters, nodes, and workloads costs, and actionable recommendations to reduce costs.

Benefits

Get a granular view of your clusters & workloads costs

Effortlessly identify potential savings

Receive actionable recommendations

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
4. Zesty Insights collects your workload usage and pattern histories
5. Real-time data is analyzed by AI algorithms
6. Zesty Insights generates actionable recommendations for potential savings

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

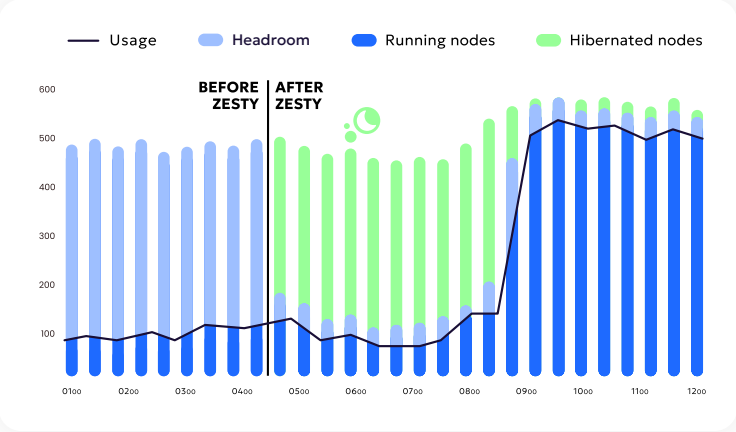

1. Zesty is granted IAM role access permission
2. Zesty installs a Kubernetes agent with a read-only-permission
3. You connect your CUR to the Zesty agent, which starts collecting workload usage and pattern histories
4. Zesty analyzes your system's scaling performance, identifies unnecessary headroom
5. 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

Spot Automation

zesty

Boost Spot instances utilization with confidence

Challenge

Spot instances utilization is risky for most applications, with AWS interrupting Spot Instances with only 2 minutes' notice, while it takes on average 5 minutes to replace an instance.

Solution

Zesty's advanced automation and our unique HiberScale technology enable Spot instance deployment in less than one minute. You can safely and automatically extend Spot instances utilization for spiky workloads, with no risk of service disruptions.

Benefits

Save up to 70%

Avoid service disruptions

Eliminate manual effort

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

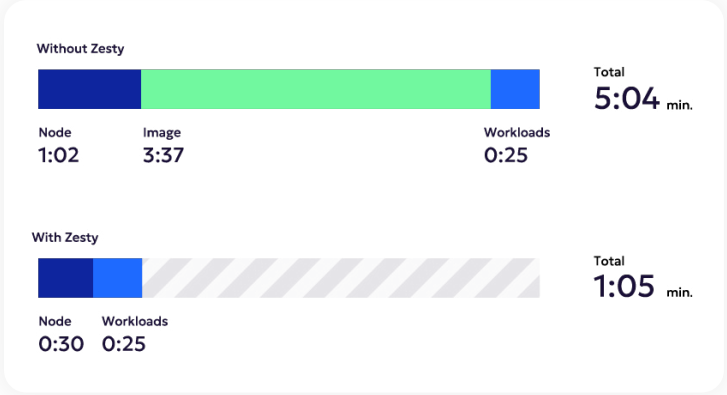

4. You install the K8s scaler, which automatically creates a pool of hibernated nodes
5. The platform identifies Spot candidates' workloads and Zesty modifies Karpenter configuration so that the identified workloads run on Spot Instances
6. When AWS sends a 2-minute notification for a Spot instance termination, a hibernated node is re-activated within 30 seconds, and a new Spot instance is scheduled before the current Spot instance terminates.
7. At any moment you can migrate back your workload to On-Demand.

Storage autoscaling

zesty

Make your PVs scalable with no risk of downtime

Challenge

Persistent volumes in Kubernetes environments lack elasticity and scalability. To prevent storage capacity failures, DevOps teams often over-provision, which results in significant unnecessary costs.

Solution

Zesty automatically scales up or down persistent volumes based on your real-time cluster needs. Our unique solution helps you save up to 70% with no risk of downtime.

Benefits

- **Reduce** storage over-provisioning
- **Cut** storage costs by up to 70%
- **Eliminate** manual operations
- **Prevent** storage capacity failures

How it works

1. Zesty creates a virtual filesystem which consists of several small storage volumes.
2. Usage metrics, instance, and disk metadata are continuously tracked and sent to Zesty's backend.
3. Zesty adds, removes, or extends persistent volumes without disrupting running pods, ensuring continuous operation and dynamic scalability.