FinOps Karpenter Implementation

Our enterprise client now uses Karpenter and AWS Spot instances for huge cloud infrastructure savings

Cost optimisation
Forecast saving of $2.4 million annually.

Zero downtime
Seamless integration with GitOps workflows and zero customer impact.

Reduced operational effort
Karpenter implementation will save 600+ engineering hours over the next year.

CHALLENGE

The client struggled with high costs and operational inefficiencies due to reliance on AWS EKS Node Groups. These rigid configurations prevented the use of cost-saving Spot Instances, requiring predefined instance types and frequent manual adjustments for changing workload needs. Scaling was slow, as EKS Node Groups depended on EC2 Auto Scaling Groups, and managing AMI updates across clusters became a burden. Engineers faced increased complexity, navigating both AWS and Kubernetes contexts, leading to higher costs, delays, and strained resources.

SOLUTION

The client replaced AWS EKS Node Groups with Karpenter, a dynamic Kubernetes autoscaler, to optimise resource usage and reduce costs. Karpenter prioritises Spot Instances for cost efficiency, dynamically provisions instances to match workloads, and automates instance lifecycle management, significantly simplifying server patching across over 100 clusters. The solution, deployed via the client's GitOps processes using Rancher and Fleet, eliminated manual intervention, enhanced agility, and reduced operational complexity, freeing up engineering resources for higher-value tasks.

2
MONTH FINOPS ENGAGEMENT

$2.4m
INFRASTRUCTURE COST REDUCTION

6000+
ENGINEER HOURS SAVED

CONCLUSION

The implementation of Karpenter significantly benefited both the client and their Platform Team. Engineers now have more time to focus on developing new features, aligning with a "Platform as a Product" mindset that adds greater value for software engineers.

Operationally, AMI upgrade efforts were reduced by an estimated 6000 hours per year, and **cloud infrastructure cost savings reached approximately $2.4 million annually across 130 clusters**. Team morale improved due to reduced operational overhead, and the platform maintained zero downtime, ensuring uninterrupted service for its users.

Thank you for reading

Do you need help with Cloud Native technologies? Get in touch and let's work together.

💬 **Talk to a Cloud Platform Consultant**

✓ Quick response   ✓ Free consultation