

» CASE STUDY

Intelligent automation enables Augury to confidently run AKS production and AI workloads on spot VMs

The challenge

Navigating complexity and streamlining VM management

By adopting cloud-native software with container architecture, Augury strategically migrated to Microsoft Azure as their primary cloud provider and adopted Azure Kubernetes Service (AKS). This move positioned them at the forefront of cloud innovation, ensuring continued operational efficiency and scalability.

Augury's AKS infrastructure was spread across three Azure accounts: development and AI, staging and production. The staging and production environments each housed a single cluster with several thousand pods, while the development and AI accounts maintained four clusters, with two dedicated to AI workloads. Initially, Augury's DevOps team handled this complex infrastructure internally. However, facing the need to reduce operational complexity, streamline cloud costs and improve gross margins, Guy Carmy, Augury's DevOps lead, sought an optimization solution for their AKS infrastructure.

"The need to cut our cloud costs became more pressing over the years," said Carmy. Prior to onboarding Ocean, Augury had been using Azure Spot Virtual Machines via Terraform. However, Azure's autoscaler proved too rigid, limiting node pools to a single virtual machine (VM) size, and yielding minimal cost savings. "We needed a more comprehensive way to manage VMs—one that could account for machine pricing and availability while delivering a more aggressive downscaling where possible."



» AT A GLANCE

Industry: Software development

Products: Flexera Ocean

Augury, a leader in machine health and process health solutions, uses purpose-built AI technology, trained by industry experts and the world's largest data library, to help manufacturing and industrial companies eliminate production downtime, improve process efficiency, maximize yield and reduce waste and emissions.

The solution

Choosing Flexera Ocean

Ocean* emerged as the ideal choice for Augury. The solution's intelligent automation dynamically scales infrastructure based on real-time analysis and workload demands of AKS clusters, continuously optimizing for performance, efficiency and cost-effectiveness. By automating key processes and providing actionable insights, Ocean empowered Augury to achieve greater operational efficiency and elevate application performance through continual optimization of cost and resource utilization.

Meeting the rigorous standards for partnership

Augury was a design partner for Ocean's AKS support, and Ocean was rigorously evaluated to meet Augury's criteria for third-party solutions:

- **Tangible and measurable value:** Ocean's ability to optimize cloud resources and streamline operations aligned perfectly with Augury's goals
- **Integration with existing workflows:** Augury values smooth integration with their existing workflows, including industry standards like Terraform and Helm charts. Ocean's compatibility with these tools ensured a smooth transition for Augury's DevOps team, minimizing disruption and maximizing efficiency
- **Compliance with CISO-defined security standards:** Augury prioritizes data security and requires any solution, including Ocean, to adhere to their CISO-defined security standards

- **Adherence to corporate procurement guidelines:** Augury insists on partnering with reputable enterprise-grade vendors. Ocean's established reputation made it an ideal choice

Ocean's enterprise-grade SaaS offering easily met these requirements, securing its position as Augury's preferred solution.

Using Azure spot VMs in production with Ocean

Following a successful Proof of Concept (POC), Augury migrated their dev, AI and staging environments to Virtual Node Groups (VNGs) managed by Ocean. Subsequently, production clusters were also migrated, achieving nearly 100% spot VM coverage—an impactful cost-saving measure.

“We quickly migrated our production to Ocean-run spot VMs,” said Carmy. “Ocean handles them efficiently, ensuring our production availability remains unaffected.”

Ocean's reliability has enabled Augury to optimize its internal developer platform. All environments autonomously spun up by Augury's developers are now hosted on Ocean-managed AKS clusters, benefiting from efficient computing and resource-effective autoscaling.

“We quickly migrated our production to Ocean-run spot VMs. Ocean handles them efficiently, ensuring our production availability remains unaffected.”

Guy Carmy
DevOps lead
Augury

The results

Revolutionizing Augury's cloud cost management strategy

With Ocean, Augury achieved significant cost savings and streamlined their AKS infrastructure through key features:

- **Bin packing:** Condensing workloads onto fewer, more utilized nodes reduced complexity and improved Augury's resource and cost efficiency
- **Proactive rightsizing suggestions:** Using Ocean's recommendations, node sizes were adjusted to actual workload demands, allowing Augury to operate more efficiently on smaller, cost-effective nodes
- **Heterogenous VNGs:** Ocean's ability to support diverse machine types within a single node pool minimized the number of pools Augury needed to manage

"Ocean's VNG flexibility delivered near-instant cost and time savings," said Carmy. "We drastically simplified our AKS infrastructure—from requiring unique Azure node pools for each workload type to consolidating multiple workloads with varying needs to run on the same Ocean VNG. Thanks to Ocean's VNGs accommodating heterogeneity, managing our infrastructure became a 'set it and forget it' process. Upon implementing Ocean, we set up two VNGs that have hosted all our workloads ever since."

Ocean's reliability ensured uninterrupted operations, supported by its seamless integration with AKS that enhances control over Kubernetes operations. Flexera empowered Augury to maximize business value in Azure operations, enabling their DevOps team to prioritize pioneering digital innovations for competitive advantage.

"It would be deeply disappointing if we lost Ocean," said Carmy. "It plays a crucial role in cost management without compromising performance or scalability. Without it, optimizing costs would be a lot more challenging, and the results might not measure."

“

If you're grappling with high AKS costs, there are two key steps to take. First, consider vertical scaling or rightsizing. This will make sure that your app only requires what it really needs. It's a tough mission at scale, so I'm really looking forward to Ocean's release of automatic rightsizing. And second, you want to onboard a cost-aware cluster autoscaler from the smarter range. Out of those, Ocean is the best possible choice we found for AKS.”

Guy Carmy
DevOps lead
Augury



Looking to the next phase of use cases for Ocean

Augury is interested in examining the development of Ocean's automatic rightsizing for AKS, a groundbreaking feature set to transform its cloud management strategy. This unique capability will facilitate continuous rightsizing in the background through automation policies, ensuring optimal resource allocation without manual intervention.

Additionally, Augury looks forward to using Ocean's support for use case-specific acceptance of rightsizing recommendations, providing flexibility tailored to their needs.

"The trust we've cultivated with Spot made us open to exploring the Spot* portfolio further," said Carmy. "Vendor consolidation saves me plenty of headache and logistics time. When I find a partner I can trust, like Spot, I want them to meet as many of my needs as possible."

* Flexera acquired Spot in March 2025.

About Flexera

Flexera helps organizations understand and maximize the value of their technology, including the rising costs and risks introduced by AI, saving billions of dollars in wasted spend. Our Flexera One platform connects the dots between what technology you have, how it is used, what it costs, and where it creates risk, helping teams take control of the increasingly complex IT estate across cloud, SaaS and on-premises. We are leading the way to unify IT asset management, FinOps and SaaS management with high fidelity data from Technopedia, our proprietary reference library of technology asset data, and intelligent automation fueled by AI. That's why thousands of global organizations rely on the Flexera One platform and Technopedia. Learn more at flexera.com

» NEXT STEPS

Run a Kubernetes workload analysis to uncover your savings opportunities with Flexera Ocean

[Get your savings analysis →](#)