

AWS Containerization With a FinOps Mindset: Achieving Efficiency and Cost Reduction

Written by:
William J. Otovic



AWS Containerization With a FinOps Mindset: Achieving Efficiency and Cost Reduction

When migrating workloads to the Amazon Web Services (AWS) Cloud, selecting the optimal configuration is paramount to controlling costs and maximizing return on investment. The federal government's Cloud Smart strategy underscores this by advocating for mission-driven cloud adoption, rather than a one-size-fits-all approach. The System Architect role involves meticulously analyzing existing environments and business requirements to craft migration strategies that include thorough Key Performance Indicators (KPIs). These KPIs are crucial for measuring progress and ensuring success where efficiency and fiscal responsibility are paramount. KPIs essential for success in the government space include:

- **Cost Reduction:** Track cost savings on IT infrastructure, software licensing and operational overhead.
- **Application Performance:** Monitor application response times, throughput and resource utilization to ensure optimal performance.
- **Security Posture:** Monitor security incidents, compliance check results and data breach rates.
- **Increased Innovation:** Measure the number of new products or services launched, as well as the time to bring new services to the market.
- **Improved Collaboration:** Track enhanced communication, project completion timelines and seamless data sharing among teams.

This paper will examine how two leading AWS container services, AWS Elastic Container Service (ECS) Fargate and Amazon Elastic Kubernetes Service (EKS), directly influence these KPIs based on organizational cloud experience within AWS, viewed through the lens of financial operations, or FinOps. While "lifting and shifting" workloads may provide a quick start to the cloud journey, it may not be the most cost-conscious option. Refactoring applications into microservices and utilizing cloud-native container services offers a significant opportunity for increased flexibility, scalability, maintainability, sustainable operational cost savings, and improved return on investment (ROI).

The Container Orchestration Dilemma: Simplicity vs. Complexity for KPI Achievement

In a rapidly changing tech landscape where efficient and scalable container orchestration is crucial, organizations face a tough choice on the AWS platform. The decision between AWS ECS Fargate and AWS EKS is centered on streamlined digital services and reductions in cost. Decisions depend on the complexity of applications and choosing the proper containerization fit, which can affect how quickly and effectively organizations can meet their KPIs.

ECS Fargate: Driving KPIs Through Operational Simplicity and Predictable Costs

ECS Fargate greatly simplifies container management, helping organizations reach their KPIs by reducing administrative work and providing predictable cost benefits.

Managed Infrastructure for Cost Saving and Application Performance

Fargate is a serverless compute engine where AWS fully manages the underlying Elastic Compute Cloud (EC2) instances, including patching, scaling the host infrastructure, and updating the operating system. This reduces operational workload, directly helping with the Cost Reduction KPI by lowering

ECS Fargate greatly simplifies container management, helping organizations reach their KPIs by reducing administrative work and providing predictable cost benefits.

labor costs for operations teams. It also maintains a consistent Application Performance KPI by hiding the complexity of node management.

Simpler Setup for Greater Innovation and Better Collaboration

With Fargate, the focus is on defining container configurations and services, without needing to manage nodes or complex container manifests. This easier setup allows teams to deploy microservices quickly, directly boosting the Increased Innovation KPI (faster time to market). The lower complexity also boosts the Improved Collaboration KPI as developers and operations teams spend less time on infrastructure details and more time on application features and backlog items.

Native AWS Integration for Security Posture and Cost Savings

Fargate seamlessly connects with other AWS services such as Amazon CloudWatch for logging and monitoring, Elastic Load Balancing for traffic distribution, and Identity and Access Management (IAM) for access control. This “out-of-the-box” integration uses AWS’ strong security features. For the Security Posture KPI, Fargate offers strong workload isolation, with each task running in its own dedicated environment. AWS manages the security of the underlying infrastructure as part of the shared responsibility model, easing the operational burden for patching and securing hosts. This reduces the need for complex setup or third-party tools, simplifying the Security Posture and helping achieve Cost Savings by lowering tooling expenses and management efforts. Teams oversee security at the container and task level using familiar AWS tools.

Lower Learning Curve for Cost Reduction and Increased Innovation

For teams new to container orchestration, Fargate provides a much easier learning curve. This enables new team members to become productive more quickly and allows existing teams to spend less time on complex set-ups and troubleshooting. As a result, there are direct cost reductions in training and operational time, while also supporting quicker iterations that lead to increased innovation.

Pay-as-You-Go Granularity for Optimal Cost Reduction

Fargate bills based on the vCPU and memory resources containers request and are running, rounded to the nearest second with a one-minute minimum. This precise billing minimizes wasted spending from over-provisioned resources, a common pitfall in traditional setups. Fargate's serverless model automatically manages compute resources, reducing idle waste and directly improving the Cost Reduction KPI. Fargate Spot Instances further enhances this by offering up to 70% discounts for fault-tolerant workloads, a clear win for Cost Reduction.

Simplified Cost Allocation for Cost Reduction

Because you're billed directly for vCPU and memory per task, it's easier to attribute costs to specific microservices or teams. This transparency is a cornerstone of FinOps, enabling better accountability and informed decisions for continuous Cost Reduction.

AWS EKS provides unmatched flexibility and control, which can be critical for certain KPIs. diligently.

Kubernetes (including EKS): Powerful Control with Higher KPI Investment

AWS EKS provides unmatched flexibility and control, which can be critical for certain KPIs. Still, it typically comes with a higher administrative burden and greater overall costs, which can impact cost reduction unless managed diligently.

Significant Operational Overhead (Higher Labor Costs) Impacting Cost Reduction

Even with managed EKS, teams remain responsible for overseeing worker nodes (EC2 instances), including right-sizing, group scaling, AMI patching, and maintenance. The inherent complexity of Kubernetes requires specialized expertise. This leads to higher staffing costs for Kubernetes engineers and longer resolution times for issues, directly counteracting the Cost Reduction KPI.

Complex Configuration and Ecosystem for Application Performance and Increased Innovation (With Overhead)

Kubernetes relies heavily on declarative YAML manifests, providing granular control over the Application Performance KPI. Its vast ecosystem (including Helm, Istio, Prometheus and Grafana) supports complex Increased Innovation use cases. However, integrating and maintaining these tools adds to the administrative workload and associated labor costs, increasing complexity for Improved Collaboration and raising overall expenses, which negatively impacts cost reduction. Debugging in this distributed environment can also be complex and time-consuming.

Shared Security Responsibility and Configuration for Security Posture

While AWS manages the EKS control plane security, teams are responsible for securing the worker nodes, container images, network policies, and access controls within the cluster. Achieving a strong Security Posture requires in-depth Kubernetes knowledge to configure Role-Based Access Control (RBAC), network policies and Pod Security Standards (or Admission), and integrate with external security tools. This level of granular control offers immense power but also greater responsibility for misconfigurations, which can potentially lead to security vulnerabilities if not managed expertly. This higher administrative burden can indirectly impact the Cost Reduction KPI due to the need for specialized security expertise and tooling.

Resource Waste and Optimization Difficulty Impacting Cost Reduction

It's common to over-allocate CPU and memory limits for pods, resulting in wasted resources. Continuously optimizing node utilization to minimize idle EC2 instances remains a challenge, directly impacting the Cost Reduction KPI. The fixed hourly fee for the EKS control plane also adds to baseline costs, accumulating even when worker nodes scale down. Optimizing reserved instances and savings plans in highly dynamic Kubernetes environments is complex and requires dedicated FinOps effort, further complicating Cost Reduction.

The overall operational and learning curve costs, which affect the Cost Reduction and Improved Collaboration KPIs, remain higher than with pure ECS Fargate.

Multi-Cloud/Hybrid Portability for Increased Innovation and Future-Proofing

For organizations prioritizing the Increased Innovation KPI through consistent deployments on premise or across different cloud providers, Kubernetes is the de facto standard. While this supports strategic agility and prevents vendor lock-in, it does not inherently contribute to direct Cost Reduction within a single cloud environment like AWS compared to Fargate.

Extreme Customization for Application Performance

Kubernetes offers unparalleled control over infrastructure, networking and scheduling, fulfilling unique Application Performance requirements that Fargate's abstractions might not cover. This level of control supports highly specific needs but often comes at the expense of simplicity and leads to increased operational costs.

EKS Fargate (Hybrid): Balancing Control and Reduced Management

Running Kubernetes pods on Fargate through EKS offers a balance by offloading worker node management. While it can reduce some compute-related costs by using Fargate's pay-as-you-go model, supporting the Cost Reduction KIP, the complexities of managing the Kubernetes control plane and understanding its declarative configurations still exist. This means that the overall operational and learning curve costs, which affect the Cost Reduction and Improved Collaboration KPIs, remain higher than with pure ECS Fargate. Regarding Security Posture, this hybrid approach

still relies on Fargate's underlying workload isolation but requires you to manage Kubernetes-native security features (including RBAC and network policies) within the EKS cluster.

Quantifiable performance information for AWS Fargate and Kubernetes is limited. Comparing the two containerization tools is challenging because they serve different purposes. Kubernetes is more of an orchestration platform that may meet business needs due to the size and complexity of applications. ECS Fargate, on the other hand, is a serverless compute engine for containers primarily designed to allow teams to run containerized applications without the provisioning overhead, ideal for less complex applications.

- **ECS Fargate** provides a simpler, hands-off approach that usually results in lower operational costs and predictable pay-as-you-go billing. Its built-in simplicity directly reduces costs by lowering engineering effort and removing infrastructure management tasks. This enables teams to focus on core application development, naturally supporting increased innovation and better collaboration by reducing friction and accelerating delivery. For security posture, it eases responsibility by handing off much of the infrastructure security to AWS.
- **Kubernetes (EKS)** offers powerful flexibility and extensive control, essential for complex or multi-cloud scenarios aiming to meet specific Application Performance or Increased Innovation KPIs. Kubernetes is a better choice

for large, complex applications that may scale to hundreds of Pods, making it suitable when you require fine-grained control over running a wide variety of workloads. However, it requires a significantly higher investment in engineering talent, specialized tools, and ongoing optimization. This often results in a higher Total Cost of Ownership for most organizations, even with managed services, due to increased labor and management complexity. While it can be finely tuned for Application Performance, achieving a similar Cost Reduction demands substantial, continuous expert effort. Additionally, it introduces complexity that can hinder Improved Collaboration if not managed with strong DevOps practices. For Security Posture, Kubernetes provides granular control but shifts more responsibility to the team for configuring and managing security policies and tools within the cluster.

In summary, the key trade-off is simplicity vs. control. Fargate is the clear winner in terms of ease of use and reduced operational overhead. However, if the application has specific requirements for hardware, cost predictability, networking, or if the organization values a multi-cloud strategy and the flexibility of the open-source ecosystem, Kubernetes offers a much more powerful and customizable solution. ■