

# EKS Auto + Capabilities vs. Fairwinds Managed Kubernetes-as-a-Service

What do AWS EKS Auto Mode and EKS Capabilities actually give you—and where would fully managed Kubernetes-as-a-Service from Fairwinds take more work off your plate?

Review this with your engineering, platform, and finance stakeholders to clarify what infrastructure work you want to keep in house and what makes sense to outsource.



LAYER	EKS AUTO + CAPABILITIES GIVE YOU...	FAIRWINDS MANAGED KAAS DOES ALL OF THIS PLUS...
<b>Cluster Infrastructure</b>	Automatic provisioning, scaling, and lifecycle management of worker nodes, with AWS managing EC2 instances, OS images, and key infra components (load balancing, IP management, block storage) for EKS Auto clusters.	Designs and implements an EKS-based architecture tailored to your security, networking, and compliance needs, then runs that infrastructure for you across AWS and other supported clouds.
<b>Node Security &amp; Patching</b>	Hardened, patched, short-lived node images with opinionated security defaults; AWS keeps the nodes it controls up to date and enforces best practices at that layer.	Owens end-to-end Kubernetes platform hardening: not just node images, but cluster configuration, add-on security, policy enforcement, and ongoing remediation for issues they detect.
<b>Core Add-ons (infra-level)</b>	Managed networking and storage add-ons (for example, CNI, load balancing, EBS CSI) within the scope of EKS Auto, with sensible defaults for many workloads.	Curates, installs, configures, and upgrades the full add-on stack you need for production (ingress, DNS, logging, metrics, backups, policy, etc.), and keeps it secure and up to date.
<b>Platform Architecture</b>	Reference architectures and best-practice guidance from AWS; your team still decides multi-cluster design, tenancy, namespaces, and how all the pieces fit together.	Takes primary responsibility for Kubernetes platform architecture and evolution—multi-cluster strategy, patterns, and changes over time based on your goals and constraints.
<b>Day-2 Operations</b>	Automation that reduces day-to-day node maintenance and some add-on toil; your team still handles cluster-level debugging, performance tuning, and most platform incidents.	A people-led service with SREs who monitor, operate, and troubleshoot your Kubernetes platform 24x7, acting as your platform team for Kubernetes.

<b>Upgrades</b>	Automated lifecycle for infrastructure in Auto's scope; you still initiate and coordinate add-on upgrades and app-level changes.	Plans and executes Kubernetes and add-on upgrades, manages API deprecations, and handles operational details so you stay current without burning internal engineering cycles.
<b>Governance &amp; Best Practices</b>	AWS shared-responsibility model: strong defaults and guardrails for what AWS manages; you remain responsible for most policies in the cluster (RBAC, quotas, pod security, cost controls).	Built-in governance and automated best-practice enforcement (via Fairwinds Insights and related tooling) for reliability, security, and cost across your Kubernetes footprint.
<b>Multi-cloud</b>	EKS-only: EKS Auto applies to AWS infrastructure and clusters.	A consistent managed Kubernetes experience across AWS and other major clouds, with one operations model and partner.
<b>Who Gets Paged</b>	Your team is still on-call for most cluster, platform, and integration issues; AWS is on-call for the underlying managed service.	Fairwinds SREs carry a big share of the pager for the Kubernetes platform, reducing internal on-call load and burnout.

## NEXT STEPS

Look across the columns above and note how much of the operational and on-call responsibility still lands on your team with EKS Auto + Capabilities versus what Fairwinds Managed KaaS handles for you. If most of the platform work and paging still sits with

you, the decision is whether you want to keep building and staffing that capability internally or hand it off. With Fairwinds, we manage the Kubernetes infrastructure and platform so your teams can focus on innovation and real business differentiation instead of keeping clusters alive.

