

Streamlining Amazon EKS Operations:

A How-To Guide



Table of contents

Introduction	03
Four Essential Pillars of Kubernetes Operations	04
Common Causes of Operational Gaps with Kubernetes at Scale	05
How to Streamline Amazon EKS Operations with Rafay's Kubernetes Operations Platform	06
Conclusion	16
Customer Experience (Case Study)	17
Appendix	18

Introduction

Enterprises around the globe have accelerated their paths to modernization with containerized applications orchestrated by Kubernetes (K8s) and the cloud. In order to function efficiently, however, containerized applications and the clusters they are built on require additional management services.

Amazon EKS leads the industry as a fully managed service designed to simplify the creation and management of Kubernetes clusters on AWS. However, as enterprises leverage the capabilities of Amazon EKS for mission critical environments, the number of Kubernetes clusters and applications in use can grow quickly. Even for customers using fully managed services like Amazon EKS, the expansion of clusters can quickly produce operational challenges that introduce hurdles to productivity, stall innovation, and compromise security.

These gaps naturally result from the complexity inherent in deploying, monitoring, securing and managing clusters across multiple data centers, zones and regions.

AWS partner Rafay Systems offers a modern operations platform that specializes in streamlining Kubernetes cluster and application lifecycle management. Built in the cloud, the platform provides Amazon EKS customers with added automation, security, visibility, and governance capabilities. Distinguished by its deep integration with Amazon EKS, Rafay layers seamlessly on top of Amazon EKS to quickly identify and fill operational gaps impacting Kubernetes clusters at scale. Rafay immediately automates critical tasks to streamline workflows for an optimized user experience.

Four Essential Pillars of Kubernetes Operations

Any enterprise leveraging a Kubernetes distribution or managed service such as Amazon EKS needs to enable four critical capabilities to create a production-grade Kubernetes operations practice. Across these areas, there are several questions that need to be answered by the enterprise directly or via the platform team. These questions become especially pertinent when supporting mission-critical applications and clusters in production. These capabilities are parts of a “Kubernetes operational gap” that must be bridged for successful Kubernetes implementations:



Automation

Key Question to Answer: How do you streamline cluster and app deployment and easily manage the lifecycle of your entire fleet?

As the number of modern applications grows so does the operational and lifecycle management burden of Kubernetes clusters and applications. Adding automation capabilities across both clusters and applications, enterprises speed the delivery and reduce the complexity of managing modern applications.



Security

Key Question to Answer: How do you ensure only the right people have access to your clusters and applications and that every action is auditable?

Mission-critical clusters and applications running in production require the highest-level of security and control to reduce risk of breaches and downtime.



Visibility

Key Question to Answer: How do you get a holistic view of the status and health of clusters and apps across your entire infrastructure?

Platform and Operations teams cannot manage or support what they cannot see. Thus, obtaining a single, enterprise-wide view of the status and health of Kubernetes clusters and applications is crucial.




Governance

Key Question to Answer: How to ensure usage of standardized clusters and application configurations and enforce them over time?

Enterprises must ensure and audit the consistency of your Kubernetes infrastructure and modern applications to comply with corporate policies and industry regulations.

Operational Gap Worsens with Kubernetes at Scale

As the number of clusters and applications across a company's infrastructure increases, the operational gap described above will grow. There are three key factors that expand the Kubernetes operation gap for enterprises:



Cluster Scale:

As the number of clusters increases, standardization, repeatability and control become increasingly complex and daunting.



Cluster Geography:

Expansion of the number of availability zones and AWS regions complicates holistic visibility and monitoring.



Cluster Usage Across Enterprise:

As multiple groups across an organization begin to leverage K8s, regulating access and permissions becomes more complex and new security and compliance risks are introduced.

How to Streamline Amazon EKS Operations with Rafay's Kubernetes Operations Platform

To bridge the Kubernetes operational gap, enterprises utilizing a K8s distribution or a managed service such as Amazon EKS should consider the four essential pillars of successful Kubernetes operations.

Rafay's Kubernetes Operations Platform (KOP) is a cloud-based platform (although an air-gapped solution also available) specifically designed for this purpose. KOP integrates with an enterprise's AWS account within minutes and streamlines Amazon EKS operations with its deep integration with AWS and Amazon EKS. The services of the platform were purpose-built to solve the common operational challenges that impede enterprises from implementing the four pillars of successful Kubernetes operations: automation, security, visibility, and governance.



Automation

When operating multiple clusters, enterprises will be challenged with managing the lifecycle of an entire fleet at once. Automation with Rafay helps enterprises streamline and standardize the management of modern applications. Automating cluster and application deployments, upgrades, and administrative tasks with Rafay also helps enterprises reduce errors, increase productivity, and deliver faster time-to-market for their modern applications.

With Rafay, enterprises can automate several repetitive tasks, including but not limited to continuous cluster and workload deployment and Kubernetes upgrades (new versions of Kubernetes and Amazon EKS are released several times a year).

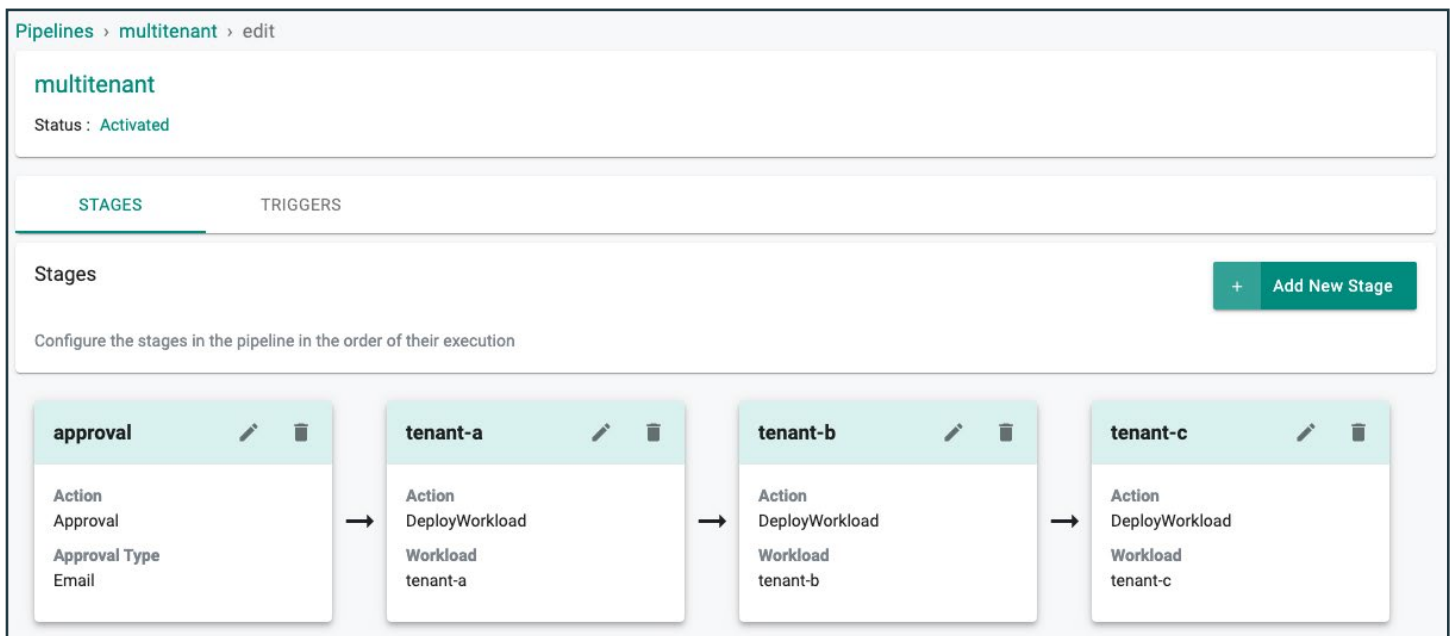
Automating cluster and application deployments

To power continuous deployment (CD), KOP includes a GitOps Service. At its core, the GitOps operating model is centered around the implementation of a version control system (e.g., Git) for a Kubernetes deployment and leveraging automation to deploy changes to clusters. GitOps is an

operating model specifically well suited for Kubernetes because it centralizes every aspect of the process for managing both operations and development.

Rafay's GitOps Service can create any number of pipelines comprised of multiple stages. The stages in the pipeline are then executed sequentially one after another.

In the example below, the pipeline has "four stages." It can either be triggered by cron job or updates to a Git repository or manually run which will send a webhook to the Controller. In this example, once the approval stage is initiated and performed by an authorized user, three workloads are deployed in a sequence.



In addition to lifecycle of k8s workloads, the Infrastructure Provisioner stage in a pipeline can be used to manage the lifecycle of infrastructure as well. As an example, [Terraform](#) by Hashicorp is supported as a "provider".

Automating Kubernetes Upgrades

For upgrading Kubernetes versions on clusters, Rafay only requires a few clicks to upgrade Amazon EKS or Amazon EKS-D clusters whether administrators require an in-place upgrade or migration to a new cluster.

KOP automatically performs all the preflight checks, upgrades the cluster, and validates results—all in matter of minutes (see image below). All of these capabilities are available through the use of a CLI and API as well.

← BACK TO UPGRADE JOBS



Target Version 1.18
 Current Version v1.17.12-eks-7684af
 Start Time 03/03/2021, 09:50:28 PM PST
 End Time 03/03/2021, 10:24:14 PM PST
 Upgrade Component Control Plane + Node Groups
 Running State **SUCCESS**
 Running Status **SUCCESS**

✓	Preflight Checks	
✓	Cluster Readiness Check	▼
✓	Nodegroup Version Check	▼
✓	Upgrade control plane	
✓	Control Plane	1602 s
✓	Upgrade Control Plane Add-ons	
✓	Add-ons	92 s
✓	Upgrade Node Groups	
✓	demo-upgrade-ng-1 SELF MANAGED	241 s
✓	Post Validations	
✓	Control Plane Version Check	▼
✓	Nodegroup Version Check	▼
✓	Nodes Ready Check	▼
✓	Pods Running Check	▼





Security

When it comes to operating multiple clusters across multiple AWS zones and regions, an organization must work a little harder to ensure that only the appropriate users are accorded permissions.

This presents some special considerations for platform teams. While most large organizations use identity management and access control for their business applications, they may group multiple clusters within a single AWS administrator account for the sake of convenience. This can create potential security vulnerabilities, and an attacker who breaches an account can also potentially gain access to all the clusters within it.

While many systems (including KOP) can define users and roles within the solution, a better approach is to use role-based access control (RBAC) governed by policies and integrated with the enterprise's preferred single sign-on solution. For example, a user in the "developer" group would have access only to clusters allocated to developers, as well as those allocated to that individual account.

Rafay's Zero-Trust Access Service was built for this purpose. The Service secures access to the Amazon EKS clusters' API server via a proxy for centralized authentication, authorization, and auditing. This proxy has out-of-the-box integrations with popular enterprise Identity Providers solutions (IdPs) such as Okta, Ping One, AzureAD, Duo SSO and others.

Given pre-built integrations, enabling this capability across an infrastructure is easy. Enterprises simply configure their IdP solution in Rafay (Okta example):

The screenshot shows the 'New IDP' configuration page in the Rafay console. The page is titled 'Identity Providers > New IDP' and has three tabs: 'IDP CONFIGURATION' (selected), 'SP CONFIGURATION', and 'METADATA CONFIGURATION'. The 'IDP CONFIGURATION' tab contains the following fields and options:

- Name ***: A text input field containing 'openc2'. A tooltip message says 'Please add a unique name'.
- IdP Type**: A dropdown menu with 'Okta' selected.
- Domain ***: A text input field containing 'openc2.io'. A tooltip message says 'Add the email domain of the organization e.g mycompany.com'.
- Encrypted SAML Assertion**: A toggle switch that is currently turned off. A tooltip message says 'If enabled, make sure to configure the IdP to encrypt the SAML Response'.
- Group Attribute Name ***: A text input field containing 'Rafay'. A tooltip message says 'Set the name of the Group Attribute Statement in SAML Assertion to map to the group with assigned roles in the console'.

At the bottom of the page, there is a '← BACK' button on the left and a 'SAVE & CONTINUE' button on the right.

Then configure the SP (organization) details:

Identity Providers > openc2

IDP CONFIGURATION **SP CONFIGURATION** METADATA CONFIGURATION

Use the information below to create an application in your identity provider (IdP)

Assertion Consumer URL (Single sign on URL)	<input type="text" value="https://console.stage.rafay.dev/auth/v1/sso/acs/3b7e2332-3bd3-48c1-a28d-d0dbe26e48e7/"/>
SP Entity ID	<input type="text" value="https://console.stage.rafay.dev/auth/v1/sso/acs/3b7e2332-3bd3-48c1-a28d-d0dbe26e48e7/"/>
NameID Format	Email Address
Group Attribute Statement Name	Rafay
Consumer Binding	HTTP-POST

Then add and configure the Rafay application within your Okta installation (via SAML protocol):

Create SAML Integration

1 General Settings **2 Configure SAML** 3 Feedback

A SAML Settings

GENERAL

Single sign on URL ?	<input type="text" value="https://console.rafay.dev/auth/v1/sso/acs/3b7e2332-3bd3-48c1-a28d-d0dbe26e48e7/"/>
<input checked="" type="checkbox"/> Use this for Recipient URL and Destination URL	
<input type="checkbox"/> Allow this app to request other SSO URLs	
Audience URI (SP Entity ID) ?	<input type="text" value="https://console.rafay.dev/auth/v1/sso/acs/3b7e2332-3bd3-48c1-a28d-d0dbe26e48e7/"/>
Default RelayState ?	<input type="text"/>
If no value is set, a blank RelayState is sent	
Name ID format ?	<input type="text" value="EmailAddress"/>
Application username ?	<input type="text" value="Okta username"/>
Update application username on	<input type="text" value="Create and update"/>

[Show Advanced Settings](#)

What does this form do?
This form generates the XML needed for the app's SAML request.

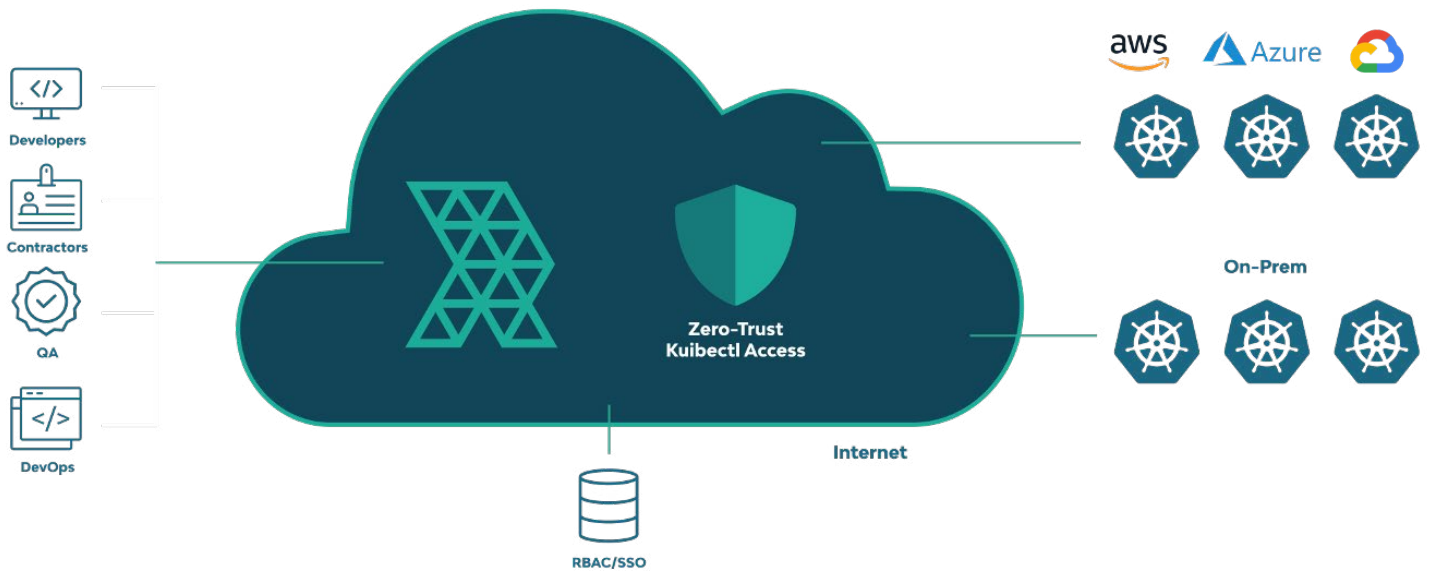
Where do I find the info this form needs?
The app you're trying to integrate with should have its own documentation on using SAML. You'll need to find that doc, and it should outline what information you need to specify in this form.

Okta Certificate
Import the Okta certificate to your Identity Provider if required.

[Download Okta Certificate](#)

More details, feature descriptions and how-to's can be found in [Rafay's product documentation](#).

The Service enables users to access clusters from anywhere—even from behind firewalls—while maintaining a full audit trail (by user and including all commands executed). Enabling this service, delivers controlled, audited access for developers, SREs and automation systems to Kubernetes infrastructure, with just-in-time service account creation and user-level credentials management, including using [AWS Identity and Access Management \(IAM\)](#). The architecture of the Service is depicted below:

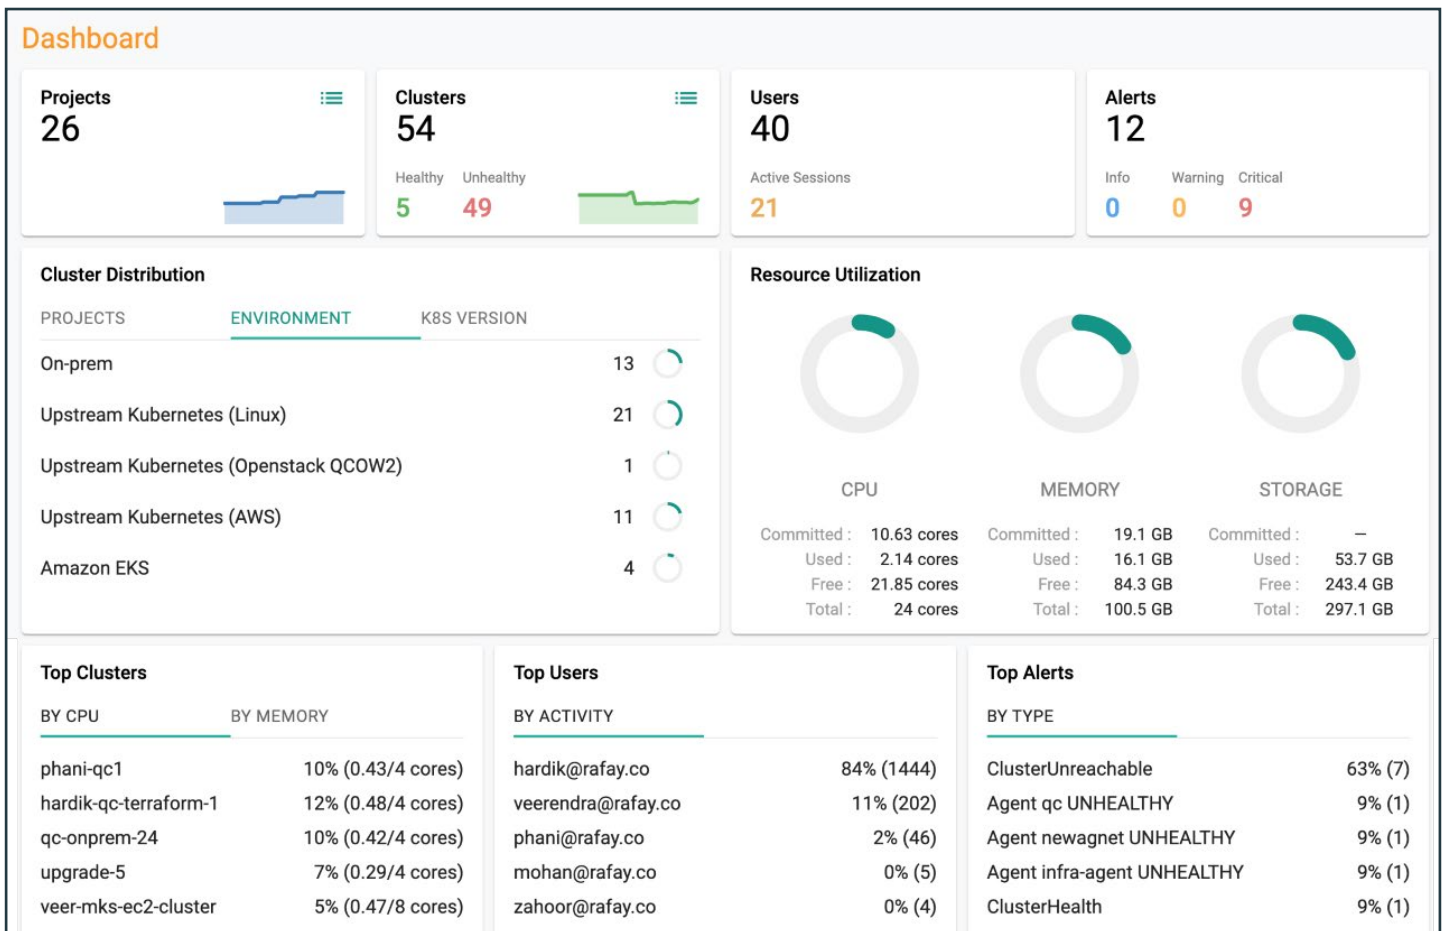




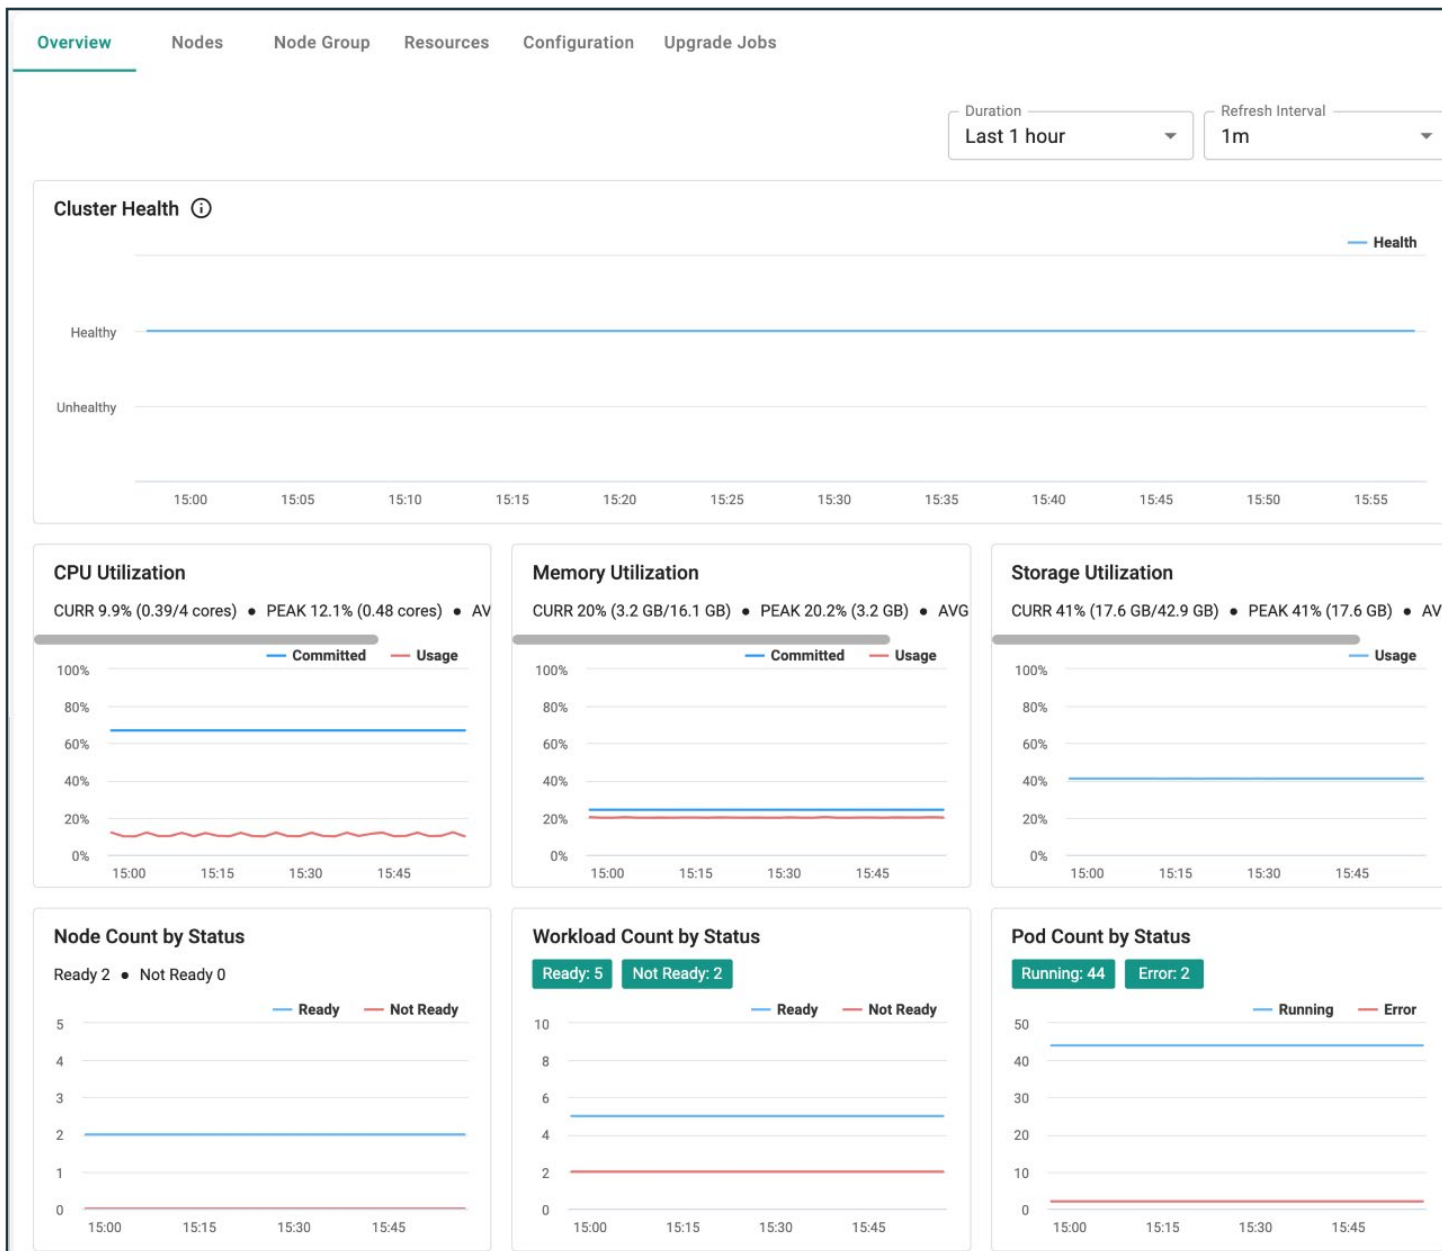
Visibility

Gaining and maintaining visibility into the health and functionality of clusters and applications is fundamental for an organization’s operations, SRE, and support teams. To ensure that resources are effectively utilized and managed across multiple clusters and AWS regions, Rafay provides enterprises with “single pane” visibility and monitoring across all infrastructure (including on-premises and remote/edge locations, no matter the K8s distribution employed).

With Rafay’s Visibility and Monitoring Service, this capability is enabled out of the box, by either importing or creating new clusters. KOP’s integrated dashboard gives cluster administrators detailed visibility and insight into the Amazon EKS (and EKS-D) clusters. Global dashboards help administrators instantly visualize, diagnose, and resolve incidents with proactive monitoring and alerting across the organization.



A detailed dashboard is available for every Amazon EKS and Amazon EKS-D cluster. It provides at-a-glance information including health, resource utilization, and pod/workload status while simplifying troubleshooting by allowing you to drill into the cluster/issue.





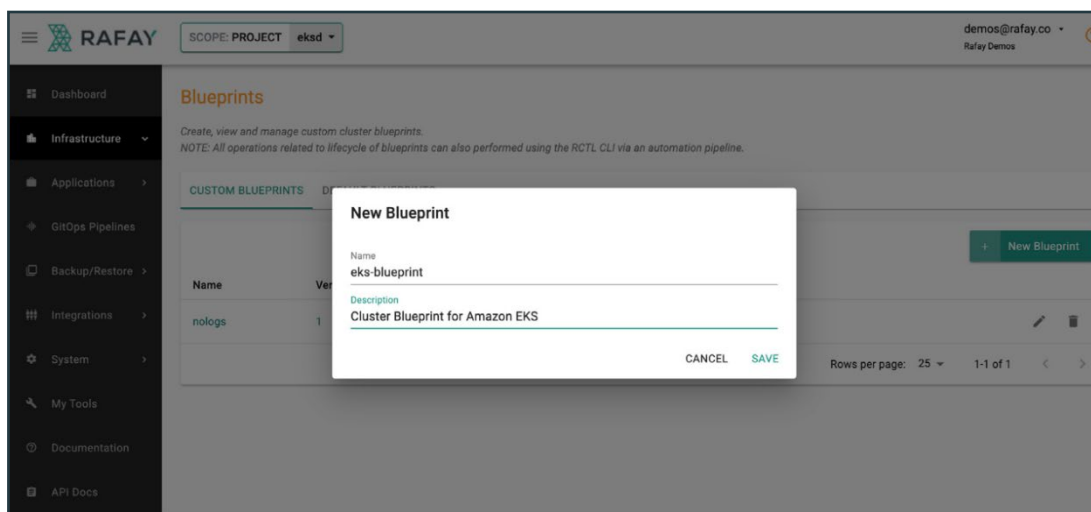
Governance

Ensuring compliance with internal policies and industry regulations (such as HIPAA, PCI or GDPR) is a fundamental requirement for a newly operational Kubernetes infrastructure. Rafay helps enterprises meet these important standards and maintain and audit compliance over time. Enterprises can leverage Rafay to quickly generate automated auditable workflows with standardized and approved templates for clusters and applications. Rafay can also quickly detect, block and notify enterprise administrators of any changes within cluster and application configurations, thus eliminating out-of-state clusters and potential security and support issues.

The best approach to governing the use of Kubernetes is to standardize configuration options for elements such as security, storage, and visibility—and apply them consistently. This vastly simplifies administration and makes patching and upgrading easier at scale. Ideally, multiple sets of pre-approved cluster configurations can be used by different internal groups and at different development stages. Rafay uses the concept of cluster blueprints to enable this capability.

With KOP, platform teams can create one or more sets of reusable cluster blueprints. Each blueprint centrally specifies configurations for clusters encompassing security policy and software add-ons such as service mesh, ingress controllers, and backup and restore solutions.

Creating a blueprint takes just a few clicks. Start by naming the new blueprint:



All KOP cluster blueprints are automatically version-controlled. Customers can also manage versions in their Git repositories. The next step is to configure the blueprint with the appropriate add-ons. See below for an example of a custom blueprint.

Blueprints > eks-blueprint > New Version

New Version Name: eks-blueprint

Create a new version of a cluster blueprint based on the configuration provided below

Version Name * *Provide a unique version for the cluster blueprint*

Pod Security Policies

Select one or more PSPs from the available list

rafay-privileged-psp x | v

PSP Policy Type ▾

Add-Ons

Select one or more addons and their versions from the available list

Name	Version	
amazon-cloudwatch	v1.24	🗑️
<input type="button" value="+ ADD DEPENDENCY"/>		
alb-ingress-controller	v1.1.8	🗑️
<input type="button" value="+ ADD DEPENDENCY"/>		
calico	v1.6	🗑️
<input type="button" value="+ ADD DEPENDENCY"/>		
nginx-ingress-controller-v035	v035	🗑️
<input type="button" value="+ ADD DEPENDENCY"/>		
<input type="button" value="+ ADD MORE"/>		

Managed System Add-Ons

The list of default addons that will be part of this cluster blueprint.

Ingress Controller	<input type="checkbox"/>
Log Aggregation	<input type="checkbox"/>
Monitoring & Alerting	<input checked="" type="checkbox"/>
Rafay K8s Operator	<input checked="" type="checkbox"/>

The entire history of cluster blueprint versions is maintained on KOP. Once a custom cluster blueprint has been created, it can be used during initial provisioning or applied to existing clusters. Blueprints can be created centrally and shared across projects ensuring the entire enterprise—across departments and across the globe—is using standardized clusters.

If a cluster in the wild changes and doesn't conform with the appropriate blueprint, Rafay's drift detection feature notifies administrators immediately. Administrators have the option to automatically block or allow changes as per their policy.

Conclusion

As enterprises grow and thrive in the cloud with Amazon EKS, encountering operational gaps within Kubernetes clusters and applications can be expected. Built in the cloud, the Rafay platform provides deep integration with Amazon EKS and fills operational gaps quickly and at scale. As a delivered service, Rafay integrates with AWS to reliably identify and solve operational problems within hours, not months. With a streamlined user experience and operational flexibility, enterprises can scale to new heights and get the most from their investment in Amazon EKS –without having to increase IT resources.



Customer Experience

Following its partnership with AWS and Rafay Systems, SonicWall achieved its infrastructure provisioning and application roll-out targets. With a 50% increase in delivery timelines rolling out their solution within three months, SonicWall optimized efficiency and exceeded customer expectations. With Rafay, SonicWall could manage the lifecycle of many more clusters using the same resources, allowing them to maintain costs, while successfully meeting targets. SonicWall's applications are now operational in six AWS Regions and Rafay's unmatched support is enabling SonicWall to serve its global customer base.

“Rafay Systems delivers a turnkey offering that automates Kubernetes cluster management and application operations at scale. The solution offers a deep integration with Amazon EKS so developers and IT users can easily bring up and manage the lifecycle of EKS clusters across AWS Regions.”

Carmen Puccio, Principal Solutions Architect at AWS

Overcome operational hurdles and accelerate your performance optimization on Amazon EKS infrastructure with Rafay Systems. [Sign up for a free trial today.](#)

Appendix:

The How-To's included in this white paper highlight a sample of the most popular ways enterprises are streamlining Amazon EKS operations. However, the Rafay KOP contains many more features that further make the scaling of Amazon EKS operations easy to manage. A summary of these features is provided in the table below:

LIFECYCLE PHASE	FEATURE	RAFAY	AMAZON EKS
Configure	EKS Cluster Provisioning <ol style="list-style-type: none"> 1. Create Cloud Credentials 2. Configure EKS Cluster 3. Provision Cluster 4. Provision EKS Control Plane 5. Provision EKS Node Group 6. Deploy Rafay Operator 7. Apply Customer Specified Blueprint 	All tasks are auto provisioned	All tasks are manually implemented
	Install Helm	Auto provisioned	Manual process
Create	Deploy the Kubernetes Dashboard	Rafay delivers an elegant multi-cluster management portal that can help you operate EKS clusters in AWS as well as Rafay KMC clusters in data centers or at the Edge	The official Kubernetes dashboard is not deployed by default.
	Deploy Microservices	Auto provisioned	Manual process
	Deploy Liveness & Readiness Probes	Auto provisioned	Manual process
	Implement Role Based Access Control	Auto provisioned	Manual process
Scale	Implement AutoScaling	Auto provisioned	Manual process
Add/Remove Node Group	Add/Remove Nodes	Auto provisioned	Manual process
Upgrade	Upgrade Nodes	Auto provisioned	Manual process
Delete	Retire Nodes	Auto provisioned	Manual process