RAFAY

How Enterprise Platform Teams Can Accelerate AI/ML Initiatives

WHITE PAPER

CONTENTS

The AI Inflection Point

Technology enabling AI has been improving over decades, continuously building new capabilities one step at a time.

The 2000's: The primary advances in this phase were using machines for analyzing data, finding patterns, generating insights, making predictions, and automating tasks at a pace and on a scale that was previously impossible.

The 2010's: The primary advances in this phase were with perception capabilities that enabled computer vision, detection, classification, voice recognition, and more.

The 2020's: While there have been multiple advances in various forms of AI, including image generators such as DALL-E, this decade has primarily been the phase of Generative AI based on the GPT-based large language models (LLMs). LLM's ability to process massive datasets has enabled them to learn the entire history, context and intent. In a nutshell, the working theory is that anything that can be conveyed through language can be addressed by Generative AI based LLMs.

Looking ahead, the inherent capabilities of GPT have been increasing dramatically every year. There has literally been an exponential increase in the number of parameters supported by each generation.

2019 – GPT-2 (1.5B parameters)

2020 – GPT-3 (175B parameters)

2023 – GPT-4 (Not disclosed yet! Trillions of parameters expected)

Given the recent enhanced power of GPT and other LLMs, it's no wonder why enterprises are so interested in leveraging this technology for their own applications. As a result, Platform teams play a key role as an accelerator to help their companies leverage the power of AI/ML in the most efficient way possible.

Why Use Kubernetes for AL/ML Applications?

Machine learning is about being able to process large data sets quickly and efficiently. Capabilities such as parallelization of jobs, data segmentation, hardware abstraction, and batch processing are critical for machine learning.

These capabilities are supported natively in Kubernetes and as a result, it is extremely well suited for machine learning. Now, let's explore the key challenges that organizations experience supporting these initiatives.

Key Challenges with AI/ML and Kubernetes

There are several key challenges that need to be solved by Platform teams before their organization can take full advantage of AI/ML applications, including:

Infra Setup and Maintenance Complexity

One of the biggest challenges organizations encounter is with the complexity of infrastructure setup and maintenance. It was recently reported that data scientists are forced to invest 60-80% of their time on infrastructure related tasks and only 4% on actual testing with data.

This is unacceptable from a user productivity perspective. Organizations want infrastructure to be abstracted away from data scientists and deliver this to them “on demand” via a “self service” experience.

Steep Learning Curve

Data scientists have a difficult enough job learning and keeping up with constant advances in the AI/ML space. It is not practical to expect them to become experts in Kubernetes and the associated ecosystem as well. There are excellent projects such as Kubeflow, mlflow that attempt to streamline this experience for data scientists. But, these still require users to be intimately familiar with Kubernetes.

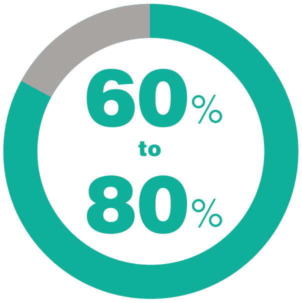

of data scientists’ time are forced to invest on infrastructure related tasks and **only 4% on actual testing with data***

Security & Governance

As AI/ML goes mainstream supporting the primary revenue stream and customer experiences for organizations, these teams find themselves having to demonstrate that they are operating with world class security and governance. Thus, defining network and security policies and proving adherence to those policies with necessary auditing become paramount.

User Access

This is an acute and daily problem where we see organizations struggle to provide data scientists and other associated users “secure, remote access” to both infrastructure and the ML system/platform.

To ensure uptime, these users need visibility into the health metrics for the underlying compute, storage infrastructure, GPUs and their applications. The lack of an integrated, intuitive access experience can result in dramatic loss of user productivity and potentially system downtime.

* https://www.mckinsey.com/capabilities/quantumblack/our-insights/rethinking-ai-talent-strategy-as-automated-machine-learning-comes-of-age

How Rafay Empowers Enterprises to Scale Their AI Initiatives

AI applications have revolutionized industries, but managing and scaling these complex applications – and the large language models (LLMs) powering them – can be a daunting task. This is where Rafay Systems steps in, offering a unified platform from which enterprises can streamline and automate the lifecycle of AI applications and their dependencies, enabling Platform teams and their companies to unlock the full potential of their AI initiatives.

Rafay's Cloud Automation Platform provides a solution to platform teams for building automated self-service cloud infrastructure workflows. The solution allows platform teams to enable anyone who depends on rapid access to cloud infrastructure to move faster safely.

With our solution, platform teams deliver **autonomy** to cloud users (developers, data scientists, researchers, and more) while maintaining **control** and **efficiency** over cloud operations.

The world of ChatGPT, OpenAI, and LLMs in AI is moving fast and it's imperative that your company leverage the benefits before your competition. Building AI-powered applications is one thing, but the infrastructure setup and maintenance of these AI applications across your infrastructure is another (that's why OpenAI runs Kubernetes). Rafay makes this easy with GPU integration, unified provisioning, lifecycle management, and monitoring of AI applications no matter where they reside.

AUTONOMY

CONTROL

EFFICIENCY

AUTONOMY

Deliver Autonomy with a Self-Service Experience for Engineers and Data Scientists

With Rafay's integration with Backstage and other tools, allow your engineers and data scientists to deploy, view, and monitor all of your AI workloads and clusters in any environment, in any region.

Rafay customers in the public cloud use Rafay Cluster Blueprints to provide their users with a self-service experience. Platform teams create validated cluster templates with the entire infrastructure stack (i.e. fully functional Kubernetes clusters preloaded with all the required software for AI/ML). Data scientists can then use these pre-validated cluster templates to provision their environments on demand.

In a nutshell, with Rafay, data scientists

- Do not require expertise with features/services in the cloud
- Do not require expertise in IaC such as terraform or GitOps
- Do not require any form of "privileged access" to cloud infrastructure to provision using the templates
- Do not need to wait for days or weeks for ephemeral infrastructure to do their job

Using pre-validated cluster templates, data scientists can literally provision complete Kubernetes based operating environments for AI/ML based on Kubernetes with a click of a button.

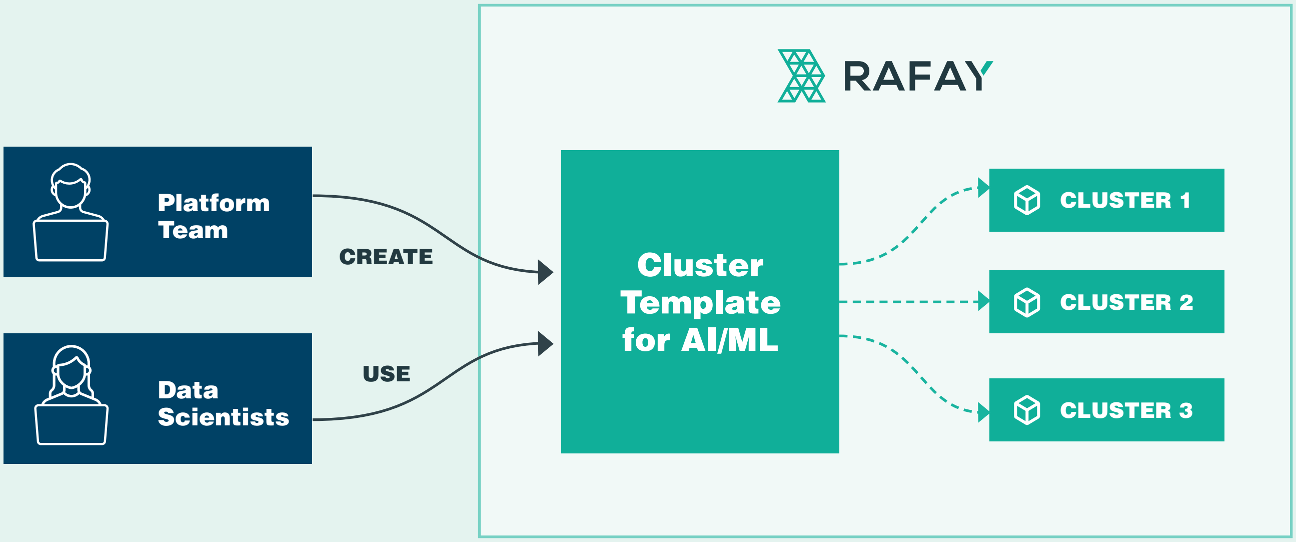

RAFAY

CONTROL

Obtain Control by Providing World-Class Standardization, Security and Governance

As AI/ML goes mainstream, Platform teams find themselves having to demonstrate that they are operating with world-class standardization, security and governance. With Rafay, enterprises enforce standards, RBAC, and have an end-to-end audit trail of all actions performed on Kubernetes clusters running LLM-based applications, for example.

For standardization, Rafay customers use Rafay Cluster Blueprints as a way to create and manage version controlled organization-wide standards for software add-ons to be deployed on their clusters.

For secure access, there are at least three classes of users that need access to support the organization's AI/ML systems and operations:

Employees

These are typically Data Scientists, Operations, FinOps and Security personnel that need "seamless, role based access" to do their jobs.

3rd Party/ISVs

Organizations frequently work with specialized 3rd party ISV software for AI/ML that needs to be deployed and operated in their infrastructure. Authorized employees from this ISV will need to be provided access so that they can "remotely" deploy, operate and manage the AI/ML application. A good example for this is a specialized AI based diagnostic application used at a hospital. The hospital will likely not be in the business of developing bespoke AI/ML applications since it is not core to their business.

Contractors

It is extremely common for organizations to leverage contractors extensively to support their AI/ML initiatives.

We see our customers using our multi-modal multi-tenancy capabilities extensively to support multiple AI/ML teams on the same Kubernetes cluster. It is incredibly common for organizations to have different teams share clusters in an effort to save costs. It is critical to make sure that doing this does not result in noisy-neighbor or security issues.

EFFICIENCY

Provide Enhanced Efficiency with Single Pane of Glass Management Across Public Clouds, Data Centers & Edge

As AI/ML environments continue to scale, the teams responsible for their operation burn cycles on executing manual processes, managing inconsistent policies, and fighting fires in order to maintain service availability. Factors like cluster drift, configuration variations, and inconsistent upgrades lead to complexity that steal the attention of operations teams from the requests from the data science teams that urgently need infrastructure for their projects. And lack of visibility into resource use by teams can be especially problematic when execution time on expensive GPU hardware is wasted.

Rafay can help you manage your entire fleet of AI/ML applications from a single pane of glass – across AWS, Azure, GCP (and others), in your on-premises data centers, and at the edge. Leverage a single, consistent GPU-specific dashboard to deploy (from NVIDIA, etc.), view and manage clusters and workloads across all your infrastructure. The benefits enjoyed by your operations teams include

- Streamlined operations due to automation of common cloud management procedures
- Reduced cloud costs with automated workflows for managing resource use
- Lower MTTR since fewer unique configs leads to faster troubleshooting

Key Features of Rafay Cloud Automation Platform for AI/ML Applications

With Rafay, you have one console to manage the operations of all your AI/ML applications (including LLMs) without having to install custom software, operational processes or dashboards. Key features and capabilities of the Rafay Cloud Automation Platform include:

Integrated GPU and Kubernetes Metrics

Rafay automatically captures and aggregates both Kubernetes and GPU metrics at the controller in a multi-tenant time series database. These metrics are then made available to users when they log in, governed by RBAC.

In a nutshell, with Rafay, users that are employees, ISVs and external contractors are provided with detailed cluster and GPU metrics just by logging in.

- No need to provide privileged, remote access to infrastructure
- No need to provide access to internally hosted monitoring applications

Here is an example of the integrated GPU metrics dashboard that users are presented with:

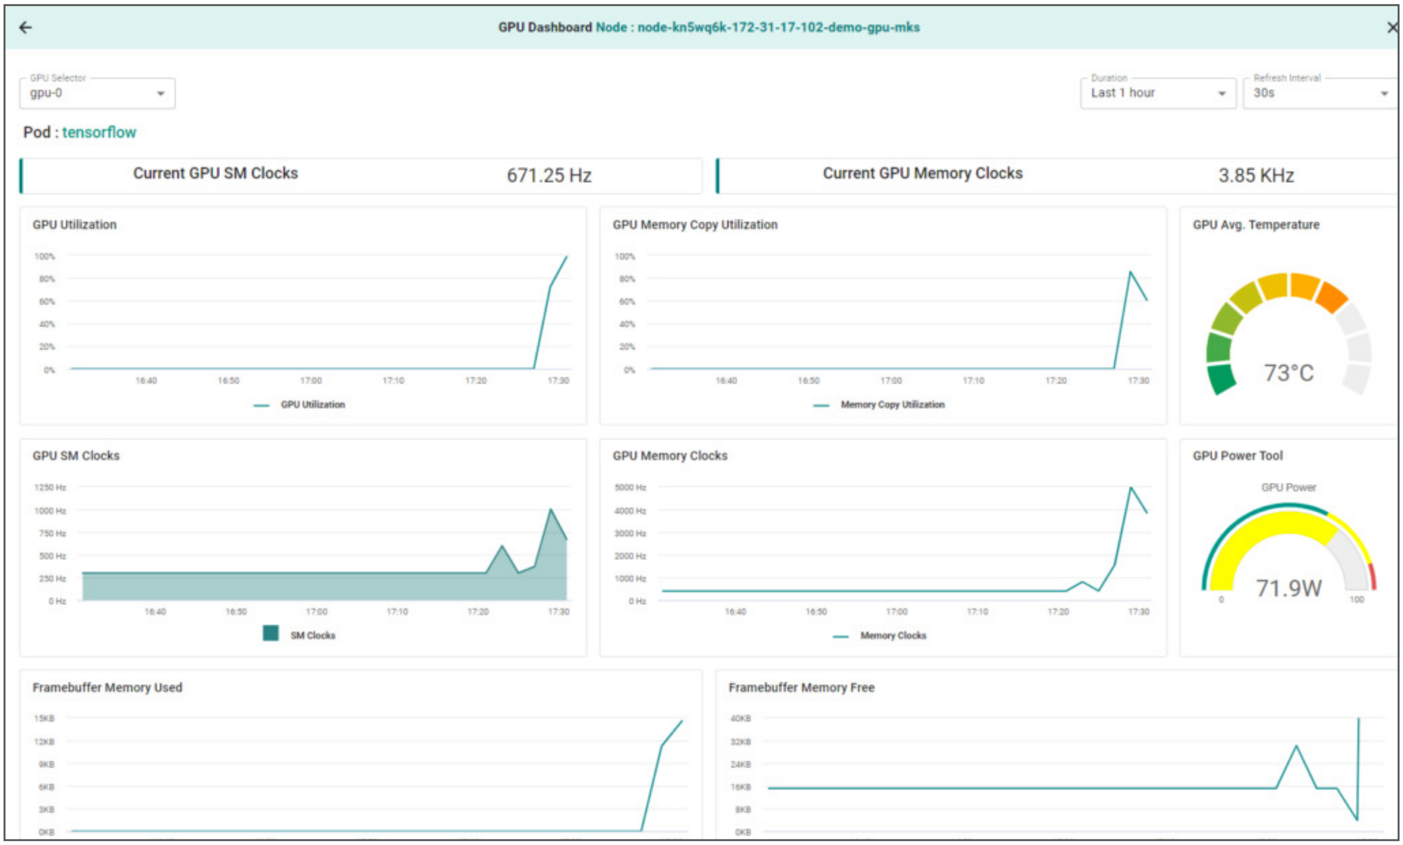

GPU Dashboard Node : node-kn5wq6k-172-31-17-102-demo-gpu-mks

GPU Selector: gpu-0

Duration: Last 1 hour

Refresh Interval: 30s

Pod : tensorflow

Current GPU SM Clocks 671.25 Hz

Current GPU Memory Clocks 3.85 KHz

Unified Management of AI/ML Apps

Organizations require a unified, central management platform for all AI/ML clusters in use spanning both data center, cloud-based and edge environments. Rafay acts as a single pane of glass to manage the deployment and lifecycle of all your AI and LLM applications.

Secure Remote Access

Users with very different roles and responsibilities (i.e. data scientists, operations, FinOps, security, contractor, 3rd party ISVs) need access and visibility into the health metrics for the underlying compute, storage infrastructure, GPUs, and their applications.

Organizations require a unified, central management platform for all Kubernetes clusters in use spanning both datacenter and cloud based environments. Rafay's unified platform acts as a proxy to govern and audit remote access.

Cluster and Workflow Standardization

Rafay's Cluster Blueprints creates and manages version-controlled standards fleet-wide for core components and software add-ons that are deployed on AI/ML clusters.

Multitenancy for AI/ML Apps

It is incredibly common for enterprises to have different teams share clusters – perhaps with specific LLM resources – in an effort to save costs. Rafay's multi-modal multi-tenancy capabilities can easily support multiple AI/ML teams on the same Kubernetes cluster.

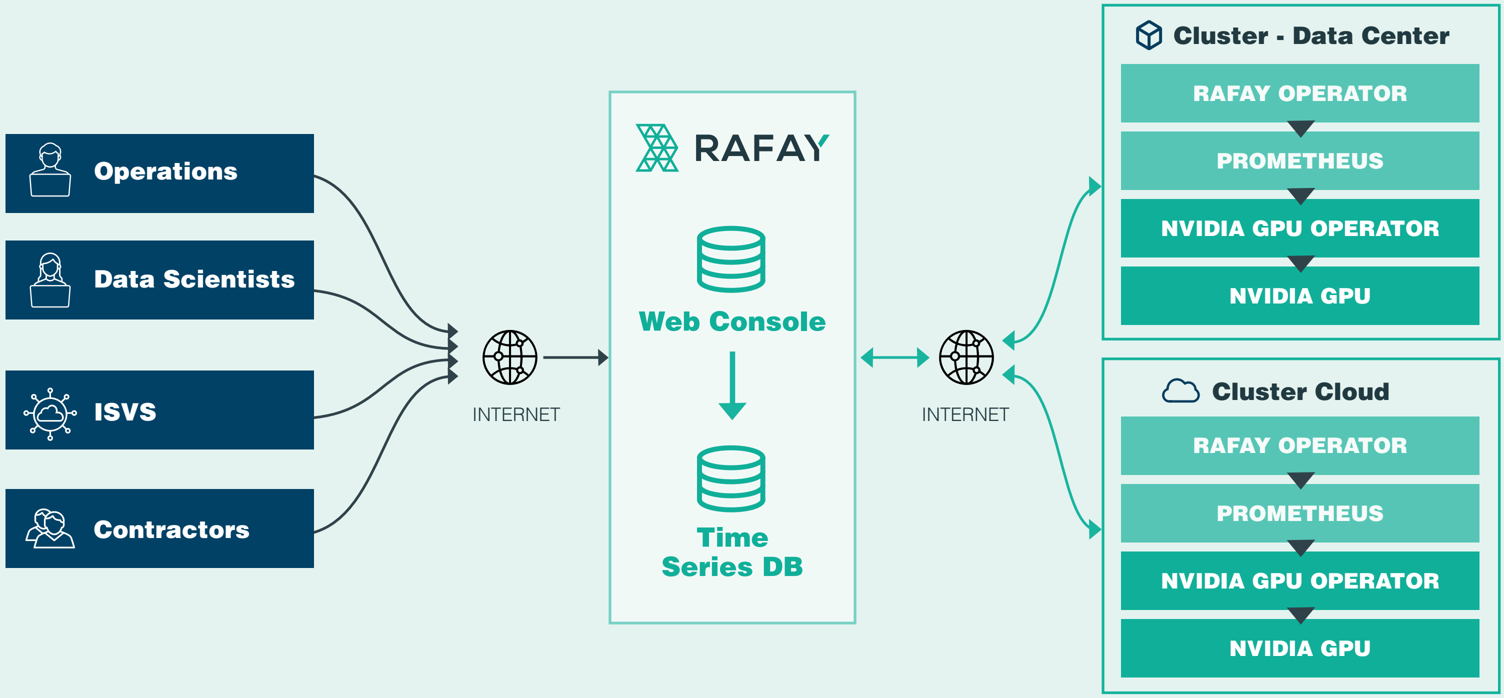

Leverage Rafay to Accelerate AI Initiatives

Rafay Systems revolutionizes AI application management by leveraging the power of Kubernetes. With integrated GPU metrics, unified management, secure remote access, cluster and workflow standardization, and multi-tenancy – all for AI/ML applications, Rafay enables Platform teams and their companies to leverage the power of AI faster, thus accelerating their AI initiatives.

Sign up for a free trial today or explore our detailed **Getting Started** guides for various AI/ML use cases.

Learn More About Rafay Systems

Rafay Systems, Inc. | 530 Lakeside Dr. Ste 220 Sunnyvale, CA 94085 | 669.336.4800 | rafay.co