

D2iQ Kaptain

An End-to-End Machine Learning Platform



Overview

Kubernetes adoption is growing at a spectacular rate as organizations improve their agility to deliver new and innovative digital services, scale to meet increasing customer demands, and achieve 99.999% uptime. With the increasing maturity of Kubernetes, its ability to scale, and improvements to the support of stateful workloads, organizations are steadily expanding the use cases for the platform.

There is a natural fit for machine learning (ML) workloads on Kubernetes, as it is well-suited to meet the scalability needs of machine learning jobs as well as embracing the continuous development nature of ML models. There are, however, several significant challenges to be overcome:

- **High risk:** Up to 87% of machine learning initiatives are abandoned before they reach production.
- **Long time to value:** For those initiatives that do make it to production, it can take more than 3 months for a single model to be deployed. Software provisioning at enterprises can take weeks or even months, which adds time and delays obtaining value.
- **New technology:** end-to-end machine learning platforms for the big data and deep learning era have only been around since 2016, with few technologies that are cloud native.
- **Complexity:** build-or-buy decisions for scalable platforms require immense knowledge of cloud-native infrastructure as well as the entire ML landscape.

Investments in ML projects that don't make it to production are completely lost. The longer it takes to get models into production, the longer it takes to get to a positive ROI (and the lower the total ROI will be).

There are also organizational challenges as multiple teams are involved in moving ML models from development to production. These include:

- Data engineers responsible for managing data acquisition and cleansing
- Data scientists responsible for creating, testing and tuning models
- Machine Learning Engineers for deploying and monitoring models
- The infrastructure teams for the platform on which all of this work needs to happen.

Key Benefits

Machine Learning on any Infrastructure

Deploy on-premise, on air-gapped networks, on the public cloud, or in a hybrid or multi-cloud environment - wherever required for mission-critical ML initiatives.

Deliver Models to Production with Speed and Agility

Break down operational barriers to seamlessly move models to production improving success of ML projects, while reducing TCO of running ML at scale.

Ready for Day 2 Operations

End to end ML requires a curated and tested set of components which extends well beyond notebooks to include integrated observability, security, and cost management.

End to End Expert Support

From basic to robust, D2iQ offers Premium and Signature support options with 24x7x365 coverage and some of the strongest SLAs in the industry to ensure that your business and machine learning initiatives have the advanced support and response time they need.

D2iQ Kaptain is an enterprise-ready end-to-end machine learning platform, powered by Kubeflow, that addresses both structural and organizational challenges and accelerates the time to market and positive ROI by breaking down the barriers between ML prototypes and production. D2iQ Kaptain enables organizations to develop and deploy machine learning workloads at scale, while satisfying the organization's security and compliance requirements, thus minimizing operational friction and meeting the needs of all the different teams involved in a successful ML effort.

Introducing D2iQ Kaptain powered by Kubeflow

D2iQ Kaptain approaches end-to-end machine learning in the enterprise “notebooks-first.” This means notebooks are the primary development environment for data science, machine learning engineering, and operations teams. We aim to give data scientists tools they are familiar with to perform tasks that are usually outside of their area of expertise: deployments and operations, as that is where the majority of enterprise efforts fail—in the last mile of ML where a positive ROI is in sight but often out of reach.

It is our belief that by making data and ML operations accessible to data scientists, our customers are more likely to succeed with their ML initiatives. This is opposed to forcing DevOps tooling, which is ill suited to the specific needs of ML, onto machine learning teams who may not be able to conform to such standards due to lacking functionality.

D2iQ Kaptain is an opinionated subset of Kubeflow projects (approx. 20+ Kubernetes Operators) with some additional components (like Spark and Horovod) that are not part of the Kubeflow ecosystem. These operators are defined with KUDO (an open source project donated by D2iQ to the CNCF) to enable full lifecycle support, which is absent from the upstream. Open-source Kubeflow does not focus on security features that enterprises require in production environments either, solved by Kaptain's end to end security model. These enhancements make Kaptain more than simply an enterprise-grade distribution of Kubeflow.

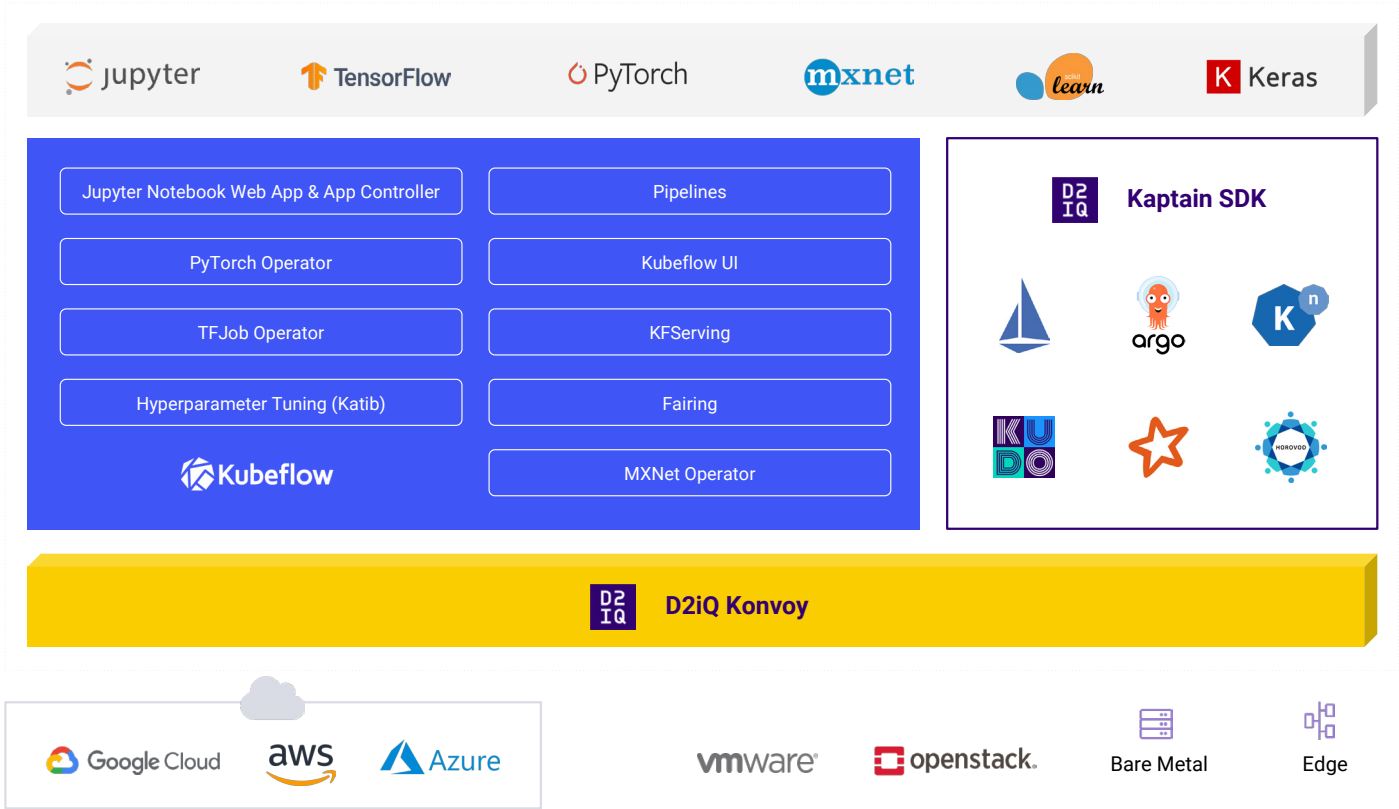
Machine Learning in Production

D2iQ ensures that investments in Machine Learning deliver ROI for your organization by ensuring that models make it to production simply. Kaptain is the enterprise machine learning platform that reduces the time needed to go from prototype to production from months to mere minutes. It bundles best-of-breed open-source technologies to build, automate, deploy, track, and monitor models across their entire lifecycle. With on-demand resource sharing, collaboration is both easy and cost effective.

World-Class 24x7 Support

D2iQ's technical support experts are masters in troubleshooting and diagnosing technologies such as Kubernetes and accompanying open source components from the CNCF ecosystem included with D2iQ Konvoy. D2iQ's best-in-class support experts have extensive experience and deep knowledge of supporting data technologies, certified operational frameworks, and machine learning components required to support enterprise deployment of machine learning initiatives in production at scale. Our full stack support offers the lowest TCO in the industry.

D2iQ Kaptain



Features	Benefits
Flexible deployment options	Deploy on any infrastructure, on-premise, public cloud, hybrid or multi-cloud environments—including highly secure air-gapped networks
Curated, integrated, and end to end tested ML platform, inclusive of popular deep learning frameworks such as TensorFlow, PyTorch and MXNet	Ensures you have all of the right components for success in your ML initiatives
Built-in and fully tested support of GPUs	Conduct intensive deep learning work without the pain of manual driver deployments
Fully automated deployment with integrated observability, end to end security, and cost management	Enterprise-grade platform ready for Day 2 operations out of the box
D2iQ Kaptain SDK	Unique to Kaptain, the SDK elevates the notebook-first approach hiding the complexities of Kubernetes and exposes what is relevant to data scientists: training, tuning, and deploying models with Python
End to end notebook tutorials	Speed up onboarding, time to value, and production readiness of ML models at scale
Compatibility with upstream open source components	Lowest TCO for end to end ML without lock-in
Cloud native with Kubernetes foundation	Modern elastic and scalable infrastructure to support standardization for both microservices and ML workloads



To learn more about about how D2iQ can be your partner in the cloud native journey, go to www.D2iQ.com.