

**D2  
IQ**

# Deliver AI/ML on Kubernetes with Speed and Agility

## Challenges

Going from prototype to production is perilous when it comes to artificial intelligence (AI) and machine learning (ML). However, many organizations struggle moving from a prototype on a single machine to a scalable, production-grade deployment. In fact, [research](#) has found that the vast majority—87%—of AI projects never make it into production. And for the few models that are ever deployed, it takes [90 days or more](#) to get there. While Kubernetes and Kubeflow would seem to be an ideal way to address some of the obstacles, the steep learning curve can introduce complexities for data scientists and data engineers who might not have the bandwidth or design to learn how to manage it.

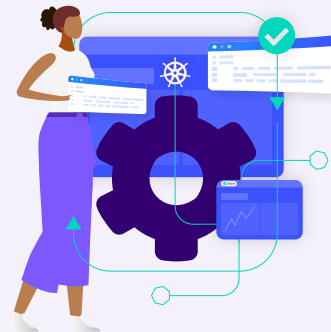


## Model Development

Jupyter Notebooks makes documentation, data visualization, and caching a lot easier for data scientists. However, as soon as you need to execute them in production at scale, notebooks become incredibly challenging to work with because they lack effective version control, testing and debugging, modularity, and extensibility. A lack of familiar tools can impede productivity and time-to-market.

## Service Complexity

Building a modern AI/ML platform from scratch requires selecting the right technologies needed for production, integrating them into the stack, and testing them to ensure they work well together and scale—all of which is time-consuming and resource intensive to set up. This leads to significant wait times for data science teams, as they, or other teams, define, build, and maintain complex environments.



## Data Access and Security

Data scientists require full access to modeling enterprise data to ensure the accuracy of their models. However, a lack of governance for deployed workloads, misconfigured access policies, and misplaced laptops can lead to significant data security risks and inefficient use of resources. As a result, data scientists are often limited to smaller, less usable snapshots of the data instead of the full data lake, which can impact model accuracy.

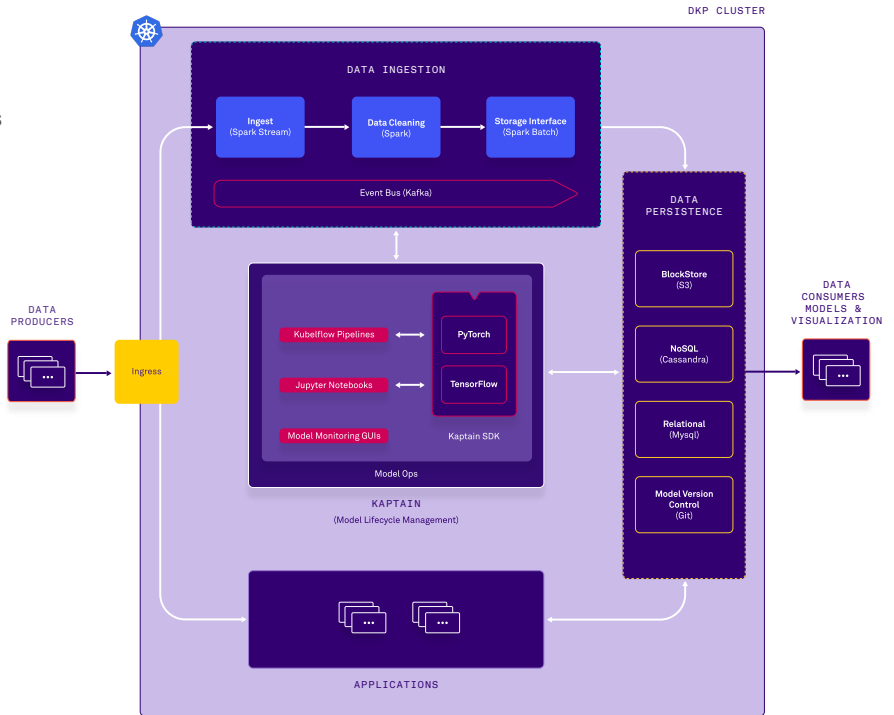
## Built for Security, Scale, and Speed

D2iQ Kaptain is an enterprise platform for AI/ML on Kubernetes for machine learning on Kubernetes that enables data scientists and data engineers to harness the scalability and flexibility of Kubernetes without having to struggle with its complexity.

D2iQ Kaptain and Kaptain SDK provides data scientists with a notebooks-first environment that is preinstalled with the best libraries, workflow tools, and frameworks, with out-of-the-box CPU or GPU support. This insulates them from the complexities of the underlying infrastructure, thus saving time and improving their productivity.

On top of providing a familiar notebook environment, Kaptain provides easy steps to create a full AI/ML workloads lifecycle, from model development to training, deployment, tuning, and maintenance. Everything can be done from within Jupyter Notebooks and the Kaptain SDK without having to switch contexts or learn Kubernetes.

Kaptain is built for success using an integrated, opinionated subset of Kubeflow and all of the necessary enterprise-grade capabilities that are validated regularly with mixed workload testing on large clusters to simulate realistic enterprise environments, at no additional cost. That way, the entire stack is guaranteed to work and scale.



In addition, D2iQ Kaptain supports fine-grained Role-Based Access Control (RBAC), as well as authentication, authorization, and end-to-end encryption with Dex and Istio.

Finally, because Kaptain and its Kubernetes underpinnings work the same everywhere, IT teams can place workloads where it makes the most sense for them.

## Outcomes

AI/ML on Kubernetes with D2iQ helps government agencies and business IT teams create unparalleled successful outcomes.

### Lower Costs

Attain positive returns on investment by reducing the complexity and friction of getting AI/ML models into production.

### Faster Deployment

Seamlessly move AI/ML prototypes to full-scale deployment with all hyperparameters tuned in a matter of minutes, not months.

### Consistency

Protect your infrastructure from malicious workloads and cryptomining attacks with a fully-integrated security stack.



D2iQ delivers the leading independent platform for enterprise-grade Kubernetes. Starting with comprehensive, enterprise-grade Kubernetes distribution built on upstream, open-source, D2iQ provides management and ancillary platform applications that are tightly integrated, secured, and tested at scale. To learn more, go to [www.D2iQ.com](http://www.D2iQ.com).