

# Success at Petabyte Scale

Excelling at Machine Learning, Artificial Intelligence,  
and Analytics on the Cloud





# The Future of Data Is Clear: It's on the Cloud

Are your machine learning (ML), artificial intelligence (AI), and analytics projects advancing at a pace acceptable to business leaders? If you rely on on-premises for big data, that platform may be holding you back. Many companies are discovering an on-premises platform has turned their data team into a bottleneck and significantly delayed their ability to execute on critical ML, AI, and analytics projects.

Companies with flourishing ML, AI, and analytics initiatives process vast amounts of data at a rapid pace and leverage different engines or frameworks based on project need. However, traditional on-premises solutions cannot handle these complexities well. Attempting to run such workloads **on-premises results in high costs, delayed insights, slow performance, and an inability to meet demands from data users.**







In contrast to on-premises solutions, a cloud-native platform like Qubole offers built-in scalability and adaptability **for faster time to value, optimized processing performance, and more sustainable economics.**

Dive into **five key benefits** achieved by switching to the cloud — and discover how a cloud-native platform can help you turn your data into profitable ML, AI, and analytics projects.

The Future of Data Is Clear:  
It's on the Cloud



# Get the Most Bang for Your Buck

When your data team conducts ML, AI, or analytics workloads, **the amount of compute required can fluctuate rapidly and send costs skyrocketing.** Unlike on-premises solutions, platforms built on the cloud decouple compute and storage — so you don't need to predict hardware requirements in advance.

**With a cloud-native platform, you save by paying only for what you use.** Qubole also provides financial governance policies and built-in workload optimization to help you save as much as 50 percent on infrastructure costs.







Qubole customers have saved up to 50 percent on CapEx and OpEx by using built-in features such as workload-aware autoscaling, aggressive downscaling, automated cluster lifecycle management, intelligent Spot management (for AWS users), and heterogeneous cluster configuration.

*"Our number one reason for choosing Qubole was we wanted to take advantage of cloud economics: only pay for what you use. Qubole's autoscaling and downscaling is definitely a huge cost saver, and the ability to isolate workloads to separate clusters is key to efficient operations."*

Oskar Austegard  
Senior Director of Data Solutions  
Gannett

Get the Most  
Bang for Your Buck





# Support More Users with Fewer Resources

A cloud platform lets organizations access virtually unlimited storage and compute on an as-needed basis. Likewise, IT teams can **scale down usage when clusters are idle or usage declines — preventing resource waste.**

Qubole automates management of the entire cluster lifecycle so that data teams can redirect their time to other important tasks. In support of different ML, AI, and analytics workloads, **Qubole automates upscaling and downscaling for any workload type and size.**







By migrating to a cloud-native platform, Qubole customers can support up to 10 times the users and data without additional administrative overhead. Qubole scales seamlessly on any cloud, enabling customers to achieve administrator-to-user ratios of 1:200 or greater.

*"The way Qubole manages the infrastructure allows me to develop in other areas of the data product lifecycle. I can spend more time on discovery, talking with clients, and really understanding their problems at a deeper level."*

Sasha Mushovic  
Data Scientist  
Return Path

Support More Users  
with Fewer Resources





# Make Data Accessible to Everyone (For Real)

A cloud-native platform supports the growing demand among users to access data without help from the data team (and in the desired timeframe and format). Qubole provides data engineers, data scientists, and data analysts with **purpose-built tools optimized for different workload types**.

Through Qubole's shared infrastructure, data users can conduct any type of big data workload leveraging Apache Spark, Hadoop/Hive, Presto, Airflow, and TensorFlow. What's more, Qubole provides partner technology, APIs, and a software development kit to deliver flexible access options for everyone.







In the last year, Qubole users processed more than 58 million commands across Apache Spark, Apache Hadoop/Hive, and Presto — an increase of 162 percent from the previous year.<sup>1</sup>

*“Our data users were changing. We needed to support not only professionals interested in descriptive reporting, but data scientists and analysts seeking to derive value from data.”*

Nathan McIntyre  
Data Engineering Manager  
Ibotta

<sup>1</sup> Qubole 2018 Big Data Activation Report

Make Data Accessible  
to Everyone (For Real)





# Rapidly Increase Productivity

Conducting ML, AI, and analytics workloads with an on-premises platform can be incredibly time-consuming and expensive. These diverse workloads **require sudden shifts in compute capacity** that can limit or disrupt productivity.

A cloud-native platform like **Qubole** makes **scaling for compute power as seamless and efficient** as possible, regardless of whether you're adding or subtracting resources, starting up/shutting down clusters, or adding/reducing nodes within a cluster.







Rapidly Increase  
Productivity

Qubole customers have reduced their data ingestion time by more than 90 percent — allowing them to better leverage critical data for near real-time dynamic insights.

*“As we made the transition to the cloud, Qubole’s ability to automate the infrastructure and automatically scale to meet the demands of our users saved us time, resources, and budget.”*

Wade Warren  
SVP of Engineering & Operations  
FANDOM (formerly Wikia)





# Maximize the Value of Your Data

Instead of spending hours (or days) combing through massive ML, AI, and analytics data sets in search of the most valuable information, a cloud-native platform optimizes performance of these tasks. Within seconds, teams are able to **rapidly discover and query data** — identifying trends and patterns that positively impact their company's bottom line.

Qubole allows users to perform all data-related tasks in the most efficient and optimal manner to **maximize the value of their data** in both the short- and long-term.







Qubole's customers can carry out ad hoc data discovery in a much more time-efficient manner to extract the greatest value from their data, even when vast amounts of information are in play. For example, one customer queries 70 million records across 300 dimensions in seconds.

*"If it wasn't for Qubole, we would have probably been delayed months to a year in embarking on our big data journey. We would have missed all the insights from the data — insights that have been a strong driver of so many of our growth strategies."*

Barkha Saxena  
VP of Data and Analytics  
Poshmark

Maximize the Value  
of Your Data







# How will the cloud benefit your ML and AI projects?

Find out what Qubole can do for you

Test Drive  
Qubole Today