



## What is Qubole?

The creators of **Facebook's Big Data infrastructure** and **founders of Apache Hive** have joined forces and leveraged their experience to deliver **Qubole** – a **Big Data cloud service offering the same advanced capabilities used by Big Data savvy organizations for a fraction of the cost.**

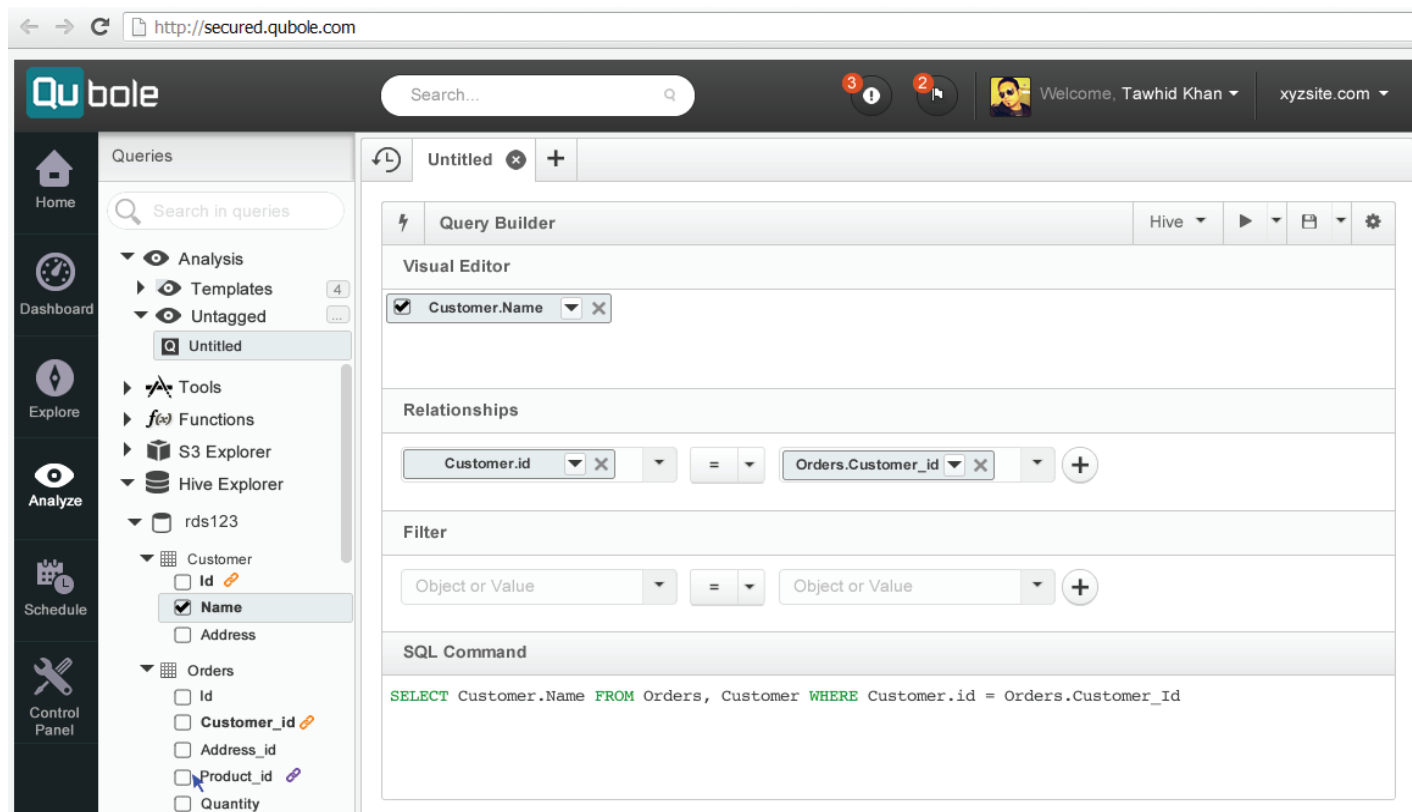


Figure1: Quboles Hive Query Builder

## Hadoop as a Service

An **elastic Hadoop cluster** in the cloud, **Hive, MapReduce, Pig** and **Sqoop** as a Service with a **Python SDK** provides all the capabilities you need to create and maintain a Big Data application.

## Improved S3 Performance

**S3 storage is secure and reliable**, but can be slow without **I/O optimization**. Qubole offers up to **5x faster query execution** against data in S3 and **2x faster data writes** than Amazon Elastic MapReduce (EMR).

## Intuitive GUI

Graphical user interface for scheduling jobs, a query editor, a visual query builder and other ways to make your job easier and more productive. Our **award winning data pipeline creator** makes it simple to setup sophisticated data ingestion, workflows and dependencies.

## Auto Scaling

Using auto scaling, Qubole **saves you money** by spinning up users' clusters only when a job is started, then automatically scales or contracts them based on the workload, and spins the servers down once the job is done.

## Optimized Hive

Our elastic cluster makes full use of daemons to **optimize resource allocation, distribution** and management. Using advanced caching and query acceleration techniques, Qubole has demonstrated query **speeds up to five times** faster than the other cloud-based Hadoop distributions.

## Spot Instance Pricing

Spot Instances allow users to **bid on unused Amazon EC2 capacity** and run those instances for as long as their bid exceeds the Spot Price. QDS makes it easy to **realize cost savings** of as much as **50% to 60%** by supporting **both the Spot and Reserved Instance pricing models.**



## Why should you care?

- **Better Utilization:** Qubole's auto scaling cluster, improved I/O optimization, and support for hybrid pricing (reserved and spot) help you realize **cost savings of as much as 50% to 60%** in total, while accomplishing tasks faster.
- **Time to Deployment:** Your managed Hadoop infrastructure is ready within minutes post signup, letting you focus on building sophisticated data pipelines, running queries, scheduling jobs and visualizing your big data
- **Availability of datasets:** Using our built-in connectors – data will start flowing into your bucket (e.g. Amazon S3), making it ready for you to cross-analyze from one location
- **Easy to Use:** manage deployment that minimizes operational interaction and provide your users with an **easy to use** solution that scales-up and down as needed without interaction with the technical team.
- **Support: 100% guarantee** that your data infrastructure is managed and supported by big data experts who have built and supported similar infrastructure for large companies (e.g. Facebook, MediaMath) and know what it takes.

## Featured Customers



## Supported Clouds



Questions? Please e-mail us at [sales@qubole.com](mailto:sales@qubole.com) or stop by our booth

[WWW.QUBOLE.COM/TRY](http://WWW.QUBOLE.COM/TRY)