



Industry: Technology

Machine Learning @Enterprise Scale

Business Problem Overview

AgilOne operates complex machine learning (ML) models and stores vast quantities of data for its customers, including major brands like Lululemon, Travelzoo, and Tumi.

AgilOne Cortex is a very robust and flexible machine learning framework built into a customer data platform. AgilOne Cortex uses supervised machine learning models to predict customer events such as purchase, subscription, and engagement. It also intelligently segments customers together based on interest and behavior using unsupervised learning techniques. AgilOne Cortex's recommender models allow the orchestration of offers and messages to customers on a 1:1 basis.

AgilOne Cortex fully supports all the functionality above with the scalability, configurability, and performance requirements needed from an enterprise machine learning platform. These features include:

- **Configurability** – AgilOne Cortex models are highly configurable to each brand's unique needs.
- **Data** – Cortex models seamlessly leverage any data sources integrated by the client, including online, offline and third party data).
- **Scale** – Cortex runs on a full cloud modular architecture that is horizontally scalable to support any models a client may envision.
- **Recipes** – Cortex provides several "recipes" of the same model, for out-of-the-box flexibility based on brand, region or any other dimensions.
- **Consumption** – Cortex model outputs are directly available in all AgilOne applications. AgilOne Actions readily exposes these model outputs thus enabling marketers to segment and personalize based on these ML model outputs seamlessly.

AgilOne operates Cortex on both Amazon Web Services (AWS) and Google Cloud Platform (GCP) and performs close to one billion predictions every day, averaging dozens of millions of customer predictions for each client across all its models.

About AgilOne

Today's enterprises, especially consumer businesses, rely heavily on customer data to acquire customers and retain their loyalty. In fact, according to Gallup, companies that leverage customer behavioral insight—what customers like, how they behave, and what they buy—report 85 percent higher sales growth than competitors that don't.

To gain real insight into customer behavior, companies must access, combine, orchestrate, and analyze a wide swath of data about existing and potential customers. This requires the use of a customer data platform (CDP). AgilOne is a leading CDP for enterprise companies. AgilOne's vision is to restore the personal relationships companies once had with customers before channel-specific marketing silos got in the way. AgilOne enables a single customer view through a best of breed identity resolution engine used on first party customer data. Machine learning and analytical algorithms enhance customer data leveraged through real-time APIs and orchestrated across all touchpoints. AgilOne helps brands build authentic omni-channel customer relationships that maximize lifetime value. The AgilOne Customer Data Platform supports more than 150 brands worldwide.

AgilOne's CDP is deployed on Google Cloud Platform and leverages Google Cloud Platform to provide unified customer data, customer intelligence and insights, and data activation and orchestration.

In order to meet the challenges of such vast amounts of data and millions of predictions, AgilOne partnered with Qubole to better automate the provision of machine learning data-processing resources based on workload, while allowing for portability across cloud providers; eliminating prototyping bottlenecks, supporting the seamless orchestration of jobs, and automating cluster management.

AgilOne now runs a variety of workloads for querying data, running ML models, orchestrating ML workflows, and more on Qubole—all on a single platform with optimized versions of Apache Spark, Apache Airflow, Zeppelin Notebooks, and leveraging Qubole's APIs to automate tasks.

Limit Bottlenecks, Simplify Cluster Management

By enabling AgilOne's data science team to make cluster management and cluster provisioning more self service, smarter, and less dependent on the Operations teams, Qubole is helping AgilOne delivery of ML models more efficiently.

Thanks to its advanced cluster lifecycle management and workload-aware auto-scaling capabilities, there's now less reliance on the Operations team as infrastructure is provisioned automatically through Qubole. Indeed, provisioning of new and larger clusters with different sets of permissions, installing dependencies on VMs, maintaining stable prototyping environments, or upgrading software is made easy. The data science team's variable infrastructure needs are now addressed by Qubole's intelligent automation—spinning up and releasing clusters and different types of nodes as needed.

Qubole's comprehensive quality assurance and support, coupled with zero-downtime software upgrades and roll-

“Qubole is strategic for AgilOne. It has helped us innovate faster on our machine learning roadmap and empower marketers with more models at scale.”

Gangadhar Konduri

Chief Product Officer, AgilOne

back capabilities, delivered the assurance and stability that AgilOne required. In Qubole's managed Zeppelin Notebooks environment, AgilOne can prototype its Python/Pyspark/Scala applications.

Cloud Portability

AgilOne supports deployment on both Amazon Web Services and Google Cloud Platform. Hence it is imperative that the underlying infrastructure used in AgilOne Cortex is cloud agnostic. With Qubole's cloud-agnostic platform, AgilOne teams can interact with data and models with a consistent user experience, build automations with the same set of APIs, and the same data-processing engines, regardless of the cloud provider.

Improved Execution

AgilOne Cortex requires a powerful orchestration system to run, monitor dozens of models for all of its clients, and to run each model across all of their users every day.

Qubole provides Airflow as a service, which allowed AgilOne's data science team, to leverage configuration-as-code workflow engine, Airflow. This allowed AgilOne to better manage the lifecycle of its ML workflows by providing easy maintenance, versioning, and testing.

Qubole also provides a comprehensive set of APIs critical for end-to-end automation. This includes automating such tasks as starting and stopping cluster usages, submitting a Spark job or changing the Spark

configuration, generating reports, increasing the timeout, and more. Execution is further enhanced by Qubole's excellent support.

"The Qubole enterprise support is of really good quality," says Arnaud Prades, Director of Data Science. "Support engineers are really knowledgeable, and escalation to higher support tiers happens faster." On a related note, setup was easy. "The setup is really easy and straightforward," says Prades. "Qubole can be up in running in about an hour or two with the right people in the room."

Looking Ahead

As its business continues to rapidly expand, the need for more data insights and more models increases. In addition to this, AgilOne is looking to leverage Qubole for running ad-hoc queries for data discovery, exploration, and analyses.

From a cluster-management perspective, AgilOne wants to further leverage Qubole's intelligent management of Google's Preemptible VMs and heterogeneous cluster management capabilities to lower its ML processing costs without compromising reliability.

"These features were among the core reasons why AgilOne partnered with Qubole," says Arnaud Prades. He also explains, "We have additional use cases for Spark, Spark Streaming and TensorFlow. In the end, we want the flexibility to use the right tool for the right job."

“ *The increased efficiency and portability across clouds are essential for AgilOne.* **”**

Gangadhar Konduri

Chief Product Officer, AgilOne

The Business Value

- Elimination of critical bottlenecks through the intelligent, autonomous, and self service provisioning of compute resources for the data science models.
- Increased efficiency for AgilOne's machine learning and operations teams.
- Improved prototyping and efficient movement of ML models into production.
- The ability to use both AWS and Google Cloud with consistent user experience, tools, and technologies.
- Efficient orchestration of machine-learning model lifecycle through Airflow.
- End-to-end automation of tasks through Qubole APIs.
- Excellent customer support, zero-downtime upgrades and roll-back capabilities.
- Fast time to value through quick and simple platform setup.

About Qubole

Qubole is revolutionizing the way companies activate their data — the process of putting data into active use across their organizations. With Qubole's cloud-native big data platform, companies activate petabytes of data exponentially faster, for everyone and any use case, while continuously lowering costs. Qubole overcomes the challenges of expanding users, use cases, and variety and volume of data while constrained by limited budgets and a global shortage of big data skills. Qubole offers the only platform that delivers freedom of choice, eliminating legacy lock in — use any engine, any tool, and any cloud to match your company's needs. Qubole investors include CRV, Harmony Partners, IVP, Lightspeed Venture Partners, Norwest Venture Partners, and Singtel Innov8. For more information visit www.qubole.com.

FOR MORE INFORMATION

Contact:
sales@qubole.com

Try Qubole for Free:
<https://www.qubole.com/products/pricing/>

469 El Camino Real, Suite 205
Santa Clara, CA 95050
(855) 423-6674 | info@qubole.com

WWW.QUBOLE.COM