



The need: computing power capable for handling high-velocity and high-volume datasets

CloudQuant is a data, research, and technology company delivering alternative solutions to financial services and other corporate clients. CloudQuant provides all tiers of users access to tools for efficient analysis of a world of fundamental, market, and consumer data through their Liberator Data Fabric and Data Warehouse services. The Chicago-based company serves a global client base.

Before becoming a standalone company in 2016, CloudQuant was an internal research product within a proprietary trading and investment firm. CloudQuant was developed to backtest and trade quantitative trading strategies algorithmically. Starting in 2017, CloudQuant's SaaS research platform became commercially available to crowd researchers globally and ultimately served researchers in 160 countries.

In operating a global research platform it became clear that the profitability of strategies developed by thousands of researchers was heavily influenced by the diversity and novelty of the datasets available to them. In order to rapidly fill this need, CloudQuant developed its flagship Liberator Data Warehouse and Data Fabric. These technology solutions depend on accessible data through web APIs and rely heavily on edge or cloud computing platforms to deliver the highest-quality service to clients. This means they need computing power capable of handling high-velocity and high-volume datasets for storage, and rapid retrieval for research as well as Extract, Transform, and Load (ETL) operations.

This translates to using cloud virtualization platforms, such as Virtual Machines (VMs) and Vultr Kubernetes Engine (VKE), to provide core services to clients, such as their CloudQuant Data Liberator and Artificial Intelligence (AI) services.

More power needed

CloudQuant needed a flexible and powerful option for additional cloud computing and cloud storage resources. Financial services data is voluminous, time-sensitive, and valuable, and CloudQuant requires reliable and easy-to-use cloud resources for clients worldwide.

”

Maintaining a high-availability and high-reliability service for large volumes of financial data can be a significant challenge, as the infrastructure needs to deliver the lowest possible latency.

Managing financial services data in the cloud is not as simple as inserting a file into a cabinet. Maintaining a high-availability and high-reliability service for large volumes of financial data can be a significant challenge, as the infrastructure needs to deliver the lowest possible latency.



<https://cloudquant.com/>

Industry

Fintech and commercial data

About CloudQuant

CloudQuant offers a high-performance quantitative research platform and data marketplace which showcases 15k+ datasets for financial services professionals.

For CloudQuant to meet these performance targets, they needed a flexible and easy-to-use service providing highly reliable cloud computing and storage resources. They achieved this through modern cloud data storage and data delivery technologies like content delivery networks that improve bandwidth and reduce access latency.

Maintaining cloud computing and storage platforms can be prohibitively expensive, so CloudQuant needed cost-effective options for the infrastructure powering their back end services, in order to ensure product line profitability. CloudQuant's strategy to achieve this was a combination of on-premises infrastructure and cloud services to minimize costs while maximizing service reliability, scalability, and flexibility. Vultr makes it easy for CloudQuant to achieve its objectives by delivering practical, easy, and cost-efficient cloud computing and cloud storage solutions.

Vultr makes it easy

CloudQuant needs to provide their clients with two primary things: robust computing power and highly reliable storage that can efficiently handle and process enormous amounts of data. Ease of use also facilitates system maintenance, and impacts long-run operating costs.

It was a significant and ongoing challenge to deliver a cloud service with the required hardware specifications and ease of maintenance to contain high-volume, high-velocity data. But with Vultr, CloudQuant's infrastructure could quickly meet their customers' architecture requirements while improving their profitability, minimizing complexity, and lowering risk.

Currently, CloudQuant uses Vultr Kubernetes Engine (VKE) and cloud computing virtual machines (VMs) with memory requirements as high as 64 GB of RAM. This ensures that their clients' compute-intensive data processing and research are able to sustain competitive performance.

With scalable cloud computing and storage provided by Vultr, CloudQuant can scale revenue by maximizing the number of services it can deliver. And developers can use open source technologies, like Kubernetes, PostgreSQL, Python, C++, and nghttp2 for their data-oriented services and applications.

Vultr's global reach – with 27 locations around the world and growing – was a key component to drive business and it has further allowed CloudQuant to meet their performance targets to improve performance and reduce latency.

An extension of the team

CloudQuant utilizes Vultr because of its high quality, low cost, and focus on customer service. They work closely with the Vultr team to customize their services. Because of the help and support they receive from Vultr, CloudQuant says they see it being a valuable partnership for years to come. Vultr's customizable approach to service enabled them to use the technologies of their choice.

”

Working with Vultr feels like we have an in-house infrastructure team at a compelling price point.

John Muehlhausen, Vice President of Engineering

As a result, CloudQuant remains exceptionally satisfied with the level of customer service they've received from Vultr. Vice President of Engineering, John Muehlhausen summed up the partnership saying, "Working with Vultr feels like we have an in-house infrastructure team at a compelling price point."