



Vertical: Technology

# Neustar Boosts Automation, Efficiency, and Savings with Qubole

## Business Problem Overview

The Neustar Unified Analytics platform helps marketers understand the impact of marketing on key business outcomes, and provides tools to enable them to optimize the allocation of their marketing investments. First, it ingests large volumes of client marketing data from a variety of sources. Then, it applies proprietary algorithms to build a predictive spend attribution model on top of that data. This reveals how the client's marketing spend correlates to revenue—enabling marketers to determine which marketing channels are working, which ones aren't, and what to do next.

This process is simple in concept but complicated in practice. The wide variety of data sources and formats ingested poses one problem. The sheer volume of data poses another. Then there's the issue of velocity—the speed at which the data is processed. Finally, there's the matter of data veracity—that is, how "clean" the data is.

To meet the demands of its growing client roster, the Neustar Unified Analytics team needed to confront the issues of variety, volume, velocity and veracity—often called the "four Vs." At the same time, the team needed to keep operational costs down. For this, it turned to Qubole.

## About Neustar Unified Analytics

Neustar Unified Analytics is an integrated marketing measurement, analytics, and attribution solution from Neustar Information Services, Inc. Neustar Unified Analytics is not a marketing campaign management tool. Rather, it runs alongside its clients' campaign management platforms to measure and attribute overall marketing spend across campaigns. More than 90 Fortune 200 companies depend on Neustar Unified Analytics to assess and improve their marketing investments.

## Addressing Variety, Volume and Velocity

Neustar Unified Analytics handles data sets from a variety of sources and in a variety of formats. This variety represents “the biggest challenge” in the data-processing pipeline, said Neustar Vice President of Systems Engineering Dan Peterson. Complicating matters are the sheer size of these data sets (50 to 60 terabytes per customer) and the volume of data processed each day (as much as 1.5 terabytes). “Just being able to process these data sets at all is challenging,” says Peterson —“let alone process them quickly”.

Neustar Unified Analytics previously used other vendor tools to process data. However, these tools just couldn't scale to meet Neustar's needs. So, the Neustar Unified Analytics team turned to Qubole. “We are able to scale up quite a lot on Qubole,” said Peterson. And “the reason why is we are able to use Spot instances a lot more with Qubole than with other platforms”, which represents anywhere from an 85–95% cost reduction. In fact, according to Peterson, at times 90% of Unified Analytics' 300–400 nodes per customer are Spot nodes. Thanks in part to Qubole's Intelligent Spot Management and comprehensive task automation, the team was able to whittle its data-processing turnaround time significantly.

## Ensuring Data Veracity

“Models are only as good as the input data,” said Peterson. “If your data has lots of gaps, then the model won't be good, no matter what algorithm you use.” But most data scientists fail to detect “dirty” data until after they run the model—a typical data science pipeline has data processing, modeling, and scoring stages. Data scientists must then fix the data and rerun the many of their processes—a task that might take weeks or even months, depending on processing speed and capacity. Often, this cycle repeats, compounding the delay. Indeed, “the main reason these things take a long time is because of the reruns,” said Peterson. This is why most organizations have difficulties with marketing campaign profitability and spend allocation: at best it's too late, but most often it's incorrect and they have no way to know it.

Neustar Unified Analytics is different. Its machine learning models include a series of pre-checks and post-checks to validate data. “Because we have a very comprehensive set of validation routines that run on Qubole, we're able to isolate problems earlier and avoid these reruns,” Peterson explains. As a result, data validation jobs require just one to one and a half run cycles. This allows Neustar to deliver insights to its clients much faster and with the highest degree of confidence.

“We work with Fortune 500 companies, so quality and performance are extremely important for us—and Qubole provides both.”

**Dan Peterson**

Neustar Vice President of  
Systems Engineering

## Keeping Costs Down

Ultimately, like anyone, the Neustar Unified Analytics team wants the highest possible performance at the lowest possible cost. This is one of the main reasons Neustar chose Qubole over other vendor tools. Aside from its Intelligent Spot Management and Workload-aware Autoscaling, Qubole's aggressive downscaling helps Neustar Unified Analytics pay only for the compute capacity, eliminating idle excess capacity.

On a given month, the heavy compute time needed for most machine learning jobs is 80 to 90 hours on average for each customer. The rest of the time is typically consumed running reports, tuning parameters, and so on—tasks that require considerably less compute power. For this reason, before Neustar Unified Analytics partnered with Qubole, its 400-odd compute nodes per customer were frequently underutilized—with no adjustment in cost. Now, Qubole aggressively—and automatically—shuts down excess capacity during slow periods, efficiently “packing” workloads in fewer nodes. This dramatically reduces operating costs, without compromising performance or delivery times.

All told, the Neustar Unified Analytics team has reduced its costs to the tune of 85 to 95 percent over its prior use of other vendor tools with reserved compute instances and administrator-led scaling. Most importantly, Qubole's cost savings don't come at the expense of customer support. “We need a very responsive support team when there are issues with the infrastructure,” says Peterson. “The support we get from Qubole is really good.”

## Looking Ahead

Looking ahead, the Neustar Unified Analytics team is focused on continuing to improve performance and its automated processes in order to deliver critical insights to clients even faster.

The flexibility, APIs, and financial governance offered by Qubole enables the Neustar Unified Analytics team to automate its solutions.

“*Qubole is cheaper and much more economical than other vendors...but more importantly, it's much more stable, and much more high-performing.*”

**Dan Peterson**

Neustar Vice President of  
Systems Engineering

“*Qubole offered us the best price for performance, and outstanding support.*”

**Dan Peterson**

Neustar Vice President of  
Systems Engineering

“*From a performance aspect, we want to be faster and faster...and Qubole fits right into this.*”

**Dan Peterson**

Neustar Vice President of  
Systems Engineering

## Business Value

- The ability to autoscale up by using Qubole's Intelligent Spot Management to decrease machine learning model turnaround from six months to three weeks, end to end
- The reduction of the cycle time required to validate model data by more than 62 percent with automation and alerts of pre- and post-model run checks
- The reduction of costs by 85 to 95 percent compared to other vendor tools through the use of Qubole's aggressive downscaling
- Unparalleled customer support from Qubole

### 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](http://www.qubole.com).

### FOR MORE INFORMATION

Contact:  
[sales@qubole.com](mailto: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](mailto:info@qubole.com)

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