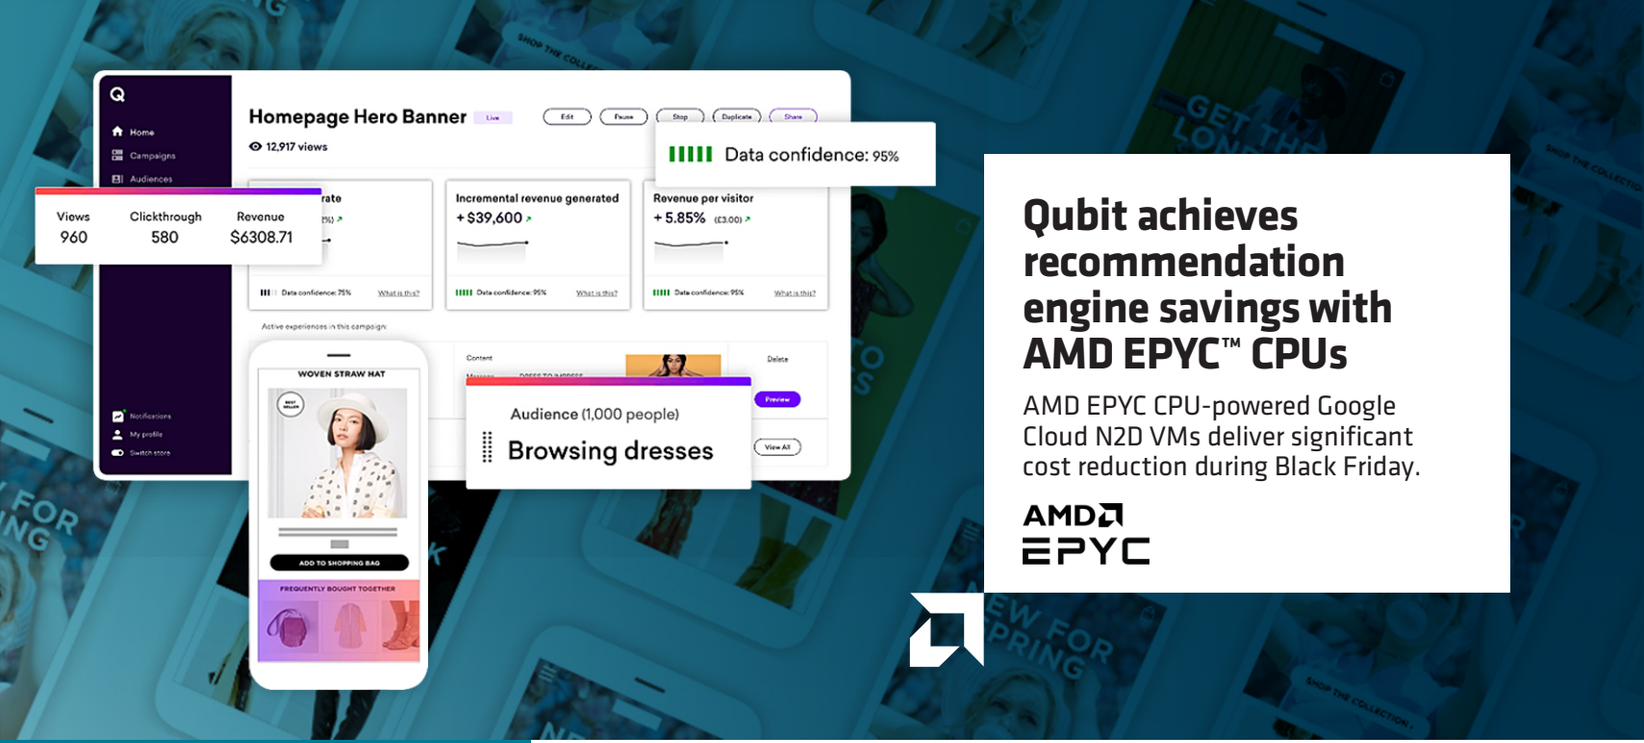

Qubit achieves recommendation engine savings with AMD EPYC™ CPUs

AMD EPYC CPU-powered Google Cloud N2D VMs deliver significant cost reduction during Black Friday.

AMD EPYC

CUSTOMER

Qubit.

INDUSTRY

Real-time personalization engine

CHALLENGES

Reduce costs while maintaining bullet proof stability

SOLUTION

Switch to AMD EPYC CPU-powered Google Cloud Platform instances

RESULTS

13 percent reduction in costs

AMD TECHNOLOGY AT A GLANCE

Google Cloud Platform N2D-highmem-4 instances powered by AMD EPYC processors

TECHNOLOGY PARTNER

Google Cloud

Product recommendations are a key part of what makes online shopping so successful. But they are most effective, if tailored to the individual shopper and delivered at just the right time. Qubit is one of the leading companies providing real-time product recommendations and insights to online retail, helping to improve customer experiences. Using its AI (Artificial Intelligence)-powered engine, Qubit delivers 9.5 million personalized experiences per hour to 364 ecommerce stores, including some of the biggest luxury brands in retail.

The recommendations provided by Qubit's engine influence 4.4 million shoppers every day and play a key role in generating $20 billion in online sales per year. For this kind of volume, reliable bullet-proof computing infrastructure is essential, particularly during peaks in traffic when online shops make the most money, such as Black Friday which is routinely the busiest U.S. retail shopping day. Qubit discovered that by switching its Google Cloud Platform provision to AMD EPYC CPU-powered instances it could deliver this demanding level of service at considerable savings.

Real-time performance is essential

"We are here to make e-commerce better," says Sergio Iacobucci, Director of Marketing and Partnerships at Qubit. "We do that through technologies built on top of Google Cloud Platform. This includes product badging within e-commerce, product recommendations, personalizing content, and changing navigations and customer journeys to align with that one person." Delivering product recommendations physically in store can be performed dynamically by the shop assistant, and email recommendations can be sent in the next newsletter. But online shopping is different. "When it comes to e-commerce, we have to react page view by page view," says Iacobucci.

"This can mean reacting within a hundred milliseconds or less. The real-time data that we collect about customers has to be queryable, fully scalable, and ready for things like Black Friday so that we can respond in the moment to the customer and personalize their experience. It is a competitive advantage of ours that we are truly real-time."

This real-time delivery needs to be always available, particularly when online shops are busy, so they can capitalize on the customer volume. "Black Friday will produce three or four times the normal traffic," says Stephen Boyle, Infrastructure Team Lead at Qubit. The company's recommendations pipeline runs on Google Kubernetes Engine (GKE). "When you go to one of our client's websites, you'll get some recommendations and they're served directly from this GKE cluster. We need to make sure this is 100 percent up all the time. We can't really afford to have things crashing."

"Running the migration across our biggest and most vulnerable time period shows that I was super comfortable with the AMD EPYC processor-powered instances."

Stephen Boyle, Infrastructure Team Lead at Qubit

Qubit wanted to maintain this reliability while optimizing costs. "I have to run all of these inexpensively as possible," says Boyle. "So it's a balance between costs, stability, and performance." The company's Google Cloud Platform representative recommended trying AMD EPYC processor-powered instances, which offered a like-for-like alternative to the N2-highmem-4 Intel-based instances Qubit had been using. The AMD EPYC-powered N2D-highmem-4 alternatives promised similar performance but with a 13 percent cost saving.

Seamless migration to AMD EPYC CPUs

Saving money would be of no use if the service provided wasn't rock solid, though. "We have really large reputable high street names as customers, and if their website goes down and we're at fault, that doesn't look good," says Boyle.

"So I need to make sure that these things run all the time, especially over peak times like Black Friday when they make most of their money. We have thousands of people shopping all the time and every one of them needs a recommendation to go with what they are buying. We have to respond in the moment, because otherwise the competitive edge of recommendations AI is removed."

"I needed to be extremely happy with the stability," continues Boyle. "So we did a lot of testing before we moved our production servers over. We're usually very willing to try any of these early access programs, testing new software or servers, because that's how you stay ahead in the game. I migrated our services live to the staging servers, so there was no downtime. We then ran them for three weeks before Black Friday to get a good feeling for how well they would perform and to see if we saw any issues. It all went really well."

This initial trial gave Qubit sufficient confidence to move its entire recommendations service over. "The week before Black Friday, we migrated our production servers," says Boyle. This consisted of 160 N2D instances with four AMD EPYC processor cores and 32GB of memory each, running GKE. "These are serving thousands and thousands of requests every minute. And there was no outage at all. We just shuffled all the traffic across. It sounds kind of boring, but nothing happened. We had no issues whatsoever. Super easy. And then the migration started saving us money!"

"I would encourage any of the Google customers who are on the fence about it to just take that leap. I would have no problems recommending the AMD service. Absolutely none."

Stephen Boyle, Infrastructure Team Lead at Qubit.

"The ten staging servers took an hour to migrate," adds Boyle. "For our production servers, we did it in about five hours, because we wanted to ease them in, just to make sure that nothing went wrong. We could have done it a lot quicker." The 13 percent cost reduction for the new AMD EPYC CPU-powered Google Cloud Platform instances was a huge saving during the high traffic peak of Black Friday. "During Black Friday, we spend a lot more money. If I can reduce that cost by 10-15 percent, that's a big win. My main goal was to try to get the new servers in before our big spikes. We caught the big holiday season spikes as well."

Trouble-free savings without compromise

"I don't know how migration could have been easier," says Boyle. "Sometimes, with CPU architecture changes, you are concerned that your code won't be quite compatible, or you'll get weird bugs. But I didn't have to worry about any performance issues or compatibility." Qubit's recommendation engine makes up most of its costs, so the overall saving was considerable too. "Our recommendation engine is about 60 percent of the workload we have," says Boyle.

"With the savings from migrating to AMD EPYC-powered Google Cloud Platform instances we can buy new cutting-edge technology and bring that to market quickly."

Sergio Iacobucci, Director of Marketing and Partnerships at Qubit

The migration was so seamless and performance so dependable, Qubit is planning to put the money saved towards new services. "With the savings from migrating to AMD EPYC CPU-powered Google Cloud Platform instances we can buy new cutting-edge technology and bring that to market quickly," says Iacobucci. "Having reduced costs in one area means we can test other stuff. That's great for us because recommendations AI is a great differentiator for us." Qubit hopes to switch its software-as-a-service workloads and machine learning to AMD EPYC processor-powered Google Cloud Platform instances in the future.

"For stability, I am 100 percent happy with how the migration went," says Boyle. "Running the migration across our biggest and most vulnerable time period shows that I was super comfortable with the AMD EPYC processor-powered instances. The collaboration between us, AMD, and Google has been excellent. I would encourage any of the Google customers who are on the fence about it, to just take that leap. I would have no problems recommending the AMD service. Absolutely none."

WANT TO LEARN HOW AMD EPYC™ PROCESSORS MIGHT WORK FOR YOU?

Sign up to receive our data center content

amd.com/epycsignup

Qubit.

RECOMMENDED JUST FOR YOU

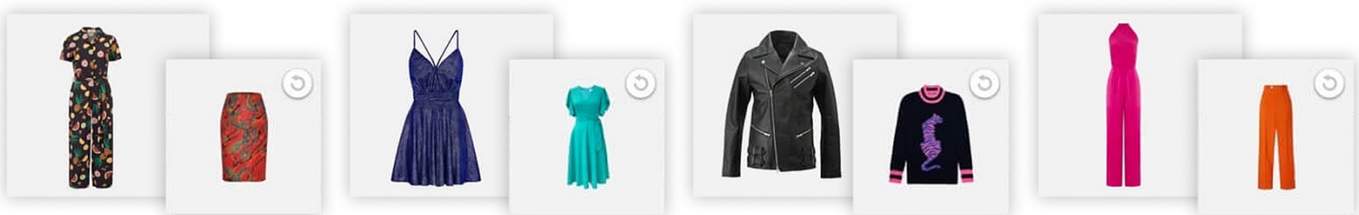

About Qubit

Qubit was founded by a small team who previously worked at Google. The company's founders believed that by using innovative techniques to collect, store and process data, they could radically transform the field of website optimization. Eight years later, Qubit's software now influences billions of visitor experiences every week. Qubit's mission is to help the world's leading retail, travel and gaming brands thrive as their customers increasingly buy online. Qubit is based in Covent Garden in the heart of London, UK. For more information visit qubit.com.

About Google Cloud

Google Cloud Platform is a suite of cloud computing services from Google for developing, deploying, and operating applications on the Internet. It is based on the same computing infrastructure as Google uses for its own internal applications. It provides a set of management tools and modular cloud services for computing, data storage, data analytics and machine learning. Google Cloud Platform earned over $13 billion in revenue in 2020, placing it in the top three cloud services globally. For more information visit cloud.google.com.

About AMD

For over 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies—the building blocks for gaming, immersive platforms, and the data center. Hundreds of millions of consumers, leading Fortune 500 businesses, and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit amd.com/EPYC.

All performance and cost savings claims are provided by Qubit and have not been independently verified by AMD. Performance and cost benefits are impacted by a variety of variables. Results herein are specific to Qubit and may not be typical. GD-181