

RBC outlines hybrid cloud strategy, eyes genAI use cases

Technological innovation doesn't always reap promised rewards. RBC's head of cloud engineering talks about why the Canadian bank's hybrid cloud strategy worked, and how it is approaching generative AI. By [Eliot Raman Jones](#)



Banks want to make money, and the argument for investing in technology is that it can make the money-making process easier. These firms may also want to be in step with the pace of tech advancement in a bid to speed up workflow processes, as well as match their more nimble, tech-forward clients.

But do these aspirations match the actual results of tech investment? A [2021 McKinsey article](#) details how often banks mismanage their technological resources and estimates that on average, banks convert just 5 to 10 cents per dollar of tech spend into additional business value. The commercial sides of banks struggle to understand the future value of emerging technologies due to the lack of immediate value, and institutions may also gamble on unknown technologies, leading to huge losses.

Vinh Tran is the head of cloud engineering and a distinguished engineer at the Royal Bank of Canada (RBC). He's been with the firm since June 2019 and views RBC as a tech company with a bank sign at the front. Tran's primary responsibility is building out RBC's cloud platforms for the more than 5,000 application teams under the bank's auspices. He believes that RBC's focus on technology and innovation sets it apart from other large banks, and points to the continued development of the bank's hybrid cloud system as an example.

"We're seeing more and more use cases for hybrid cloud now that our cloud program has matured and we embrace hybrid cloud," Tran says. "We're finding more and more opportunities to integrate public and private cloud."

waterstechnology

RBC's embrace of the cloud comes at a time when many banks still hold reservations about cloud's security and reliability for critical systems and data. Tran says the initial drive to the cloud was primarily motivated by the scaling that the public cloud offered, as the bank was unable to achieve the necessary scale for managing data on-premises.

However, while the public cloud did allow Tran's team greater freedom, he realized there were multiple use cases in which limiting the bank to operating solely within the public cloud was too big of a risk.

His team identified that the process of re-hosting—where an exact copy of an application or workload, together with its data store and operating system, is moved to the cloud—would be more expensive when applied to legacy systems operating on virtual machines (VMs) and containers.

"If you're running a legacy or traditional system that is running well on VMs or containers and you don't have a specific need for innovation or scale, there's not a big driver to move to the public cloud," Tran says. "We own the datacenters and already have these capabilities on-premises. We've been doing it for 150 years ourselves, so let's keep it there."

Stress-testing the silver lining

The movement of material and data into the cloud is a laborious process. Concerns over data security dominate discussions of public cloud implementation in banking, and Tran is acutely aware of how much of his day-to-day life is dominated by allaying data protection concerns.

"Security has always been the biggest capability that we needed to prove," he says. "It's probably 80% of my job."

RBC delayed moving data onto their public cloud until Tran's team was confident they had secured the right capabilities in the public cloud, and controlled the end-to-end experience of the data, from internal processes to the cloud platform. It took the bank several years to validate its data processes as Tran's team remained vigilant about the potential risks of public cloud adoption.

Once non-critical data had been vetted sufficiently, RBC began moving significant systems onto the cloud, but strictly in situations where the bank had identified potential value in exporting those systems. During the period of market volatility brought on by the GameStop meme stock incident, RBC's direct investing business migrated some of its mobile services to the cloud because of significant workload spikes that could not be handled effectively on-premises. However, this was not just a matter of moving the individual APIs for the services. RBC fully re-architected the service into a microservices operation.

Tran says that over the last three years, the amount of data the bank has moved to the public cloud has doubled every year. It's also a heavy user of Amazon Web Services' Kubernetes, which allows the cloud team to make software portable securely and cost-effectively.

"In our public cloud, we actually seldom use virtual machines to host applications," Tran says. "We don't think virtual machines are the future—we think containers are. We've made huge investments in container orchestration with Kubernetes."

Alex Wolcough, CEO of consultancy firm GreenBirch Group, says Kubernetes' ability to exist across multiple clouds means that software can be moved between cloud and on-premises without having to significantly recompile things.

"You can put Kubernetes in any machine you like—you can put it in the machine under your desk if you want. As long as it can see the inputs and the outputs it's expecting, Kubernetes is basically isolating the software from the physical or virtual environment it is running in. It makes the programs highly portable," Wolcough says.

Wolcough says that while the hybrid cloud model in banks is not perfect, it does allow a vast swath of applications to be moved to the cloud—but not everything.

"As far as front-office applications go, it's very unlikely they are putting anything that does high-volume, low-latency trading applications in that hybrid environment," he says. "Anything after the actual trade is done—like settlement and risk and anything that doesn't need to be done to the nearest microsecond—this hybrid environment is going to be good for."



"We're seeing more and more use cases for hybrid cloud now that our cloud program has matured and we embrace hybrid cloud ... We're finding more and more opportunities to integrate public and private cloud"
Vinh Tran, RBC

Up at night

Alongside the growth of RBC's cloud model, eyes at the executive level within RBC have been raking over the potential use cases for other emerging technologies within the bank, namely generative AI. With companies like OpenAI rolling out ChatGPT to the public as well as cloud provider-specific cases such as AWS' Bedrock service, implementing generative AI within the bank's internal structure is an appealing option for RBC.

Tran notes that the bank's data and analytics team have worked extensively with machine learning and artificial intelligence, and before the team jumps to implementing generative AI into customer-facing systems, his team's initial plans consist of adapting the technology for internal use. He cites other developer-focused programs such as Microsoft's Copilot and AWS' CodeWhisperer as examples that help developers improve their code—an area in which he says the bank has room to improve.

"Our generative AI use cases are very internally focused. We think genAI is going to be a big enabler for us, behind the scenes, making developers more productive," Tran says. He says that keeping up with the speed and innovation of technology while maintaining security of data is the main challenge for RBC going forward.

"Especially with what's happening in genAI today, we are seeing innovation and change happen at warp speed," Tran says. "That's what keeps me up at night: How do we move faster and keep up with the change while still maintaining our very high bar for security and resiliency?" **wt**