Akamai Customer Story

# GoVeda Accelerated AI-Powered Patent Intelligence

Global AI-powered patent search platform gained performance, flexibility, and long-term cloud partnership

Improved performance by up to 30%

Reduced costs by 20%

Gained strategic partnership

## Accelerating innovation through AI-powered patent intelligence

Launched in 2025, GoVeda is an AI-native patent intelligence platform that helps patent offices, legal professionals, and innovators navigate the end-to-end intellectual property lifecycle. To deliver fast, accurate search across more than 220 million patent publications worldwide, the company needed infrastructure capable of supporting highly specialized AI and vector database workloads. After a successful proof of concept, GoVeda migrated a critical search workload to Akamai Cloud, improving performance, lowering costs, and establishing a strategic partnership designed to support future growth.

## Solving one of innovation's biggest bottlenecks

For Cheng Tai, Co-Founder and CEO of GoVeda, the company's mission grew directly from personal experience. After building and exiting his first AI start-up, Tai encountered firsthand the challenges of protecting and commercializing intellectual property.

"When I filed my first patent in the United States, I learned how slow the process is," said Tai. "AI moves too fast for innovators to wait years for their ideas to be protected."

GoVeda was created to address that challenge. The Singapore-based company serves three primary audiences: patent examiners, intellectual property professionals, and inventors. Its platform combines AI-powered search, analytics, and reasoning engines to help users discover prior art, assess innovation landscapes, monitor competitors, and accelerate patent examination.

The company's long-term vision is ambitious. "We want to make the entire IP system far more efficient," explained Tai. "Our goal is to help reduce examination times from years to months while improving the commercialization and discoverability of intellectual property."

GoVeda

**Location**
Singapore
goveda.com

**Industry**
Business Services

**Solution**
Cloud Computing

**Product**
Akamai Inference Cloud

Akamai Cloud

## Building an AI platform capable of accurately searching 220 million patents

Traditional search engines aren't up to the task of accelerating patent intelligence — especially across the global database of more than 220 million patent publications that GoVeda maintains. Unlike conventional web searches that simply surface a handful of relevant results, patent examination requires exhaustive analysis.

"Examiners can review thousands of documents during a search," said Tai. "Patent documents are also written in legalese, so general-purpose AI models often struggle to understand them."

To address this challenge, GoVeda developed specialized AI models optimized specifically for intellectual property workflows. Accuracy is critical because retrieval performance deteriorates rapidly as patent databases grow larger.

"Accuracy is what ultimately makes the search process faster," explained Tai. "We benchmark extensively and find that many general-purpose systems have error rates significantly higher than our specialized models."

Supporting these AI-powered workflows requires high-performance infrastructure, particularly for vector databases that power semantic search and AI inference.

## Seeking more than another cloud vendor

Initially, GoVeda relied heavily on Google Cloud Platform (GCP) for web services, search infrastructure, and vector database workloads. As the patent search company continued to grow, however, Tai began looking beyond pure infrastructure consumption.

"We want long-term partnerships with organizations that can evolve alongside us," he said. The introduction to Akamai came through a chance conversation at an industry event. What began as a discussion about a potential collaboration quickly evolved into a technical evaluation.

Encouraged by the discussion, GoVeda's engineering team launched a proof of concept focused on several workloads, including model training and vector databases. The test delivered exactly what GoVeda hoped to see.

According to Tai, GoVeda's vector database workload needs large RAM nodes to achieve the highest performance. "Those nodes are very expensive on other cloud vendors. With Akamai, we saw we could achieve the same performance at a lower cost," recalled Tai.

But pricing was only part of the equation. Unlike hyperscale cloud providers, Akamai offered a more accessible and responsive experience.

"With hyperscalers, you're often competing with much larger customers for attention and resources," said Tai. "With Akamai, we can have in-depth conversations about how to best use the infrastructure and optimize our workloads."

**With Akamai, performance improved while costs dropped — and we're just getting started.**

**Cheng Tai**
**Co-Founder and CEO, GoVeda**

## Migrating a mission-critical search workload

Following the successful proof of concept, GoVeda migrated its vector database environment to Akamai Cloud. The workload was particularly important because it powers one of the most critical components of the user experience: search responsiveness.

Every patent search depends on rapid retrieval from large-scale vector databases. Even small improvements in response times can significantly improve productivity for patent examiners and researchers.

GoVeda's engineering team simultaneously optimized portions of its infrastructure during the migration. The result exceeded expectations.

"Our vector database performance is now 20% to 30% better than our GCP equivalent because of the migration and optimization work we completed," said Tai.

The company also realized approximately 20% savings during the proof-of-concept phase. As workloads continue to scale, Tai believes the financial benefits could become even more significant.

"As we grow and migrate more workloads, we expect the savings opportunity to increase," he explained.

## Supporting performance, scalability, and future AI growth

Today, Akamai powers GoVeda's vector database environment while providing the flexibility needed to support future expansion. For GoVeda, one of the most valuable aspects of Akamai is access to infrastructure configurations that are optimized for demanding AI workloads.

"Our search platform requires specialized vector database nodes with large memory footprints and specific RAM-to-CPU ratios," said Tai. "With Akamai, we can access those resources."

The Akamai services that are supporting GoVeda include:

- **Akamai Inference Cloud:** Infrastructure optimized for AI inference and machine learning workloads
- **CPU:** High-performance G8 dedicated 128x32 resources for AI model development and future training initiatives
- **LKE: Managed Kubernetes Engine:** Managed Kubernetes service for deploying and scaling cloud native applications
- **Akamai Object Storage:** Durable, scalable storage for large datasets and application assets
- **Akamai Block Storage:** High-performance persistent storage for critical workloads
- **Akamai NodeBalancers:** Load balancing that improves application availability and performance
- **Managed Databases:** Simplified database operations with built-in scalability and reliability

"Akamai's infrastructure just works," said Tai. "It's stable, reliable, and allows us to scale as our needs evolve."

## Preparing for the next generation of AI-driven innovation

GoVeda's traffic is already growing rapidly. Following a major industry conference in London, platform use nearly doubled. As adoption further increases, the company expects to migrate additional workloads to Akamai.

The partnership is particularly important because it's difficult to predict future AI infrastructure requirements. "It's very hard to predict workloads one or two years out," said Tai. "Flexibility is critical at our stage."

Tai believes opportunities will expand beyond the current vector database deployment. While the company continues to use Akamai's inference capabilities and cloud services, GoVeda is especially interested in Akamai's GPUs for large-scale model training initiatives.

As GoVeda prepares for continued growth, Tai sees Akamai as a partner that can help the company adapt to rapidly evolving AI demands. That approach aligns closely with GoVeda's strategy of building trusted relationships across the intellectual property ecosystem.

"I've worked with AWS, Azure, and GCP. We don't just want to buy compute resources," Tai concluded. "We want partners who help us innovate, adapt, and grow. That's what we've found with Akamai."

GoVeda

GoVeda, the AI-native platform that transforms the entire IP lifecycle and accelerates global innovation.


| Published 06/26