Akamai Customer Story

spiderSilk Scaled Global Cyber-AI with Akamai

Dubai-based cybersecurity innovator cost-effectively orchestrated AI workloads and delivered real-time threat intelligence at scale

Reduced infrastructure costs 30%–40%

Cut processing time to minutes

Reclaimed 20%–30% engineering time

Innovating global cybersecurity intelligence

Dubai-based cybersecurity innovator spiderSilk builds autonomous, AI-powered **threat detection** and exposure management solutions. As the company expanded from its 2019 startup roots to its 2025 acquisition by CPX Holding, it required cloud infrastructure that could scale in step with evolving demands. It addressed its needs by leveraging **Akamai Cloud** and globally distributed compute, Kubernetes orchestration, and edge native architecture. The result? spiderSilk is able to deliver real-time, high-throughput threat intelligence globally while maintaining predictable costs, operational efficiency, and flexibility for future growth.

Scaling AI cybersecurity without hyperscaler constraints

spiderSilk's vision has always centered on autonomous, **cloud native** cybersecurity intelligence capable of global, real-time exposure detection. Its platform ingests telemetry from around the world to identify emerging threats, then adapts AI-driven models for proactive mitigation. This approach required infrastructure that could handle burst-intensive workloads, orchestrate ephemeral compute clusters, and maintain low-latency performance for multi-region threat simulations.

During its early growth, spiderSilk considered hyperscaler solutions for certain workloads. However, these options introduced several challenges: opaque pricing models, unpredictable egress costs, limited geographic diversity for edge-based testing, and slower support resolution. Scaling workloads required extensive planning and manual intervention, which constrained rapid experimentation and deployment.

"These operational gaps highlighted a misalignment with our needs for infrastructure to support agile, burst-intensive workloads," explained Furqan Shah, Senior DevOps Lead for spiderSilk.

>spiderSilk;

Location
Dubai, United Arab Emirates
spidersilk.com

Industry
Software and SaaS

Solution
Cloud Computing

Akamai

Scaling infrastructure alongside growth with Akamai

With Akamai, spiderSilk found a cloud partner that scaled in step with its evolving requirements, offering transparent pricing, developer-centric tooling, and globally distributed infrastructure for AI-driven workloads.

In its early years, spiderSilk relied on Linode virtual machines (VMs). These VMs provided stable, reliable compute resources while giving the team flexibility to scale as workloads grew.

As AI-driven assessments became more complex, spiderSilk adopted Kubernetes-orchestrated ephemeral clusters on LKE: Managed Kubernetes Engine. This shift allowed the company to provision compute resources on demand, automatically scale during high-volume operations, and decommission resources immediately upon task completion.

"The infrastructure scaled seamlessly with us," said Shah. "From initial deployment to full multi-region operations, Akamai adapted to our needs without us having to compromise on speed or performance."

Ensuring high throughput for large-scale AI security workloads

spiderSilk's AI models rely on Akamai's infrastructure for high-throughput telemetry ingestion, inference pipelines, and anomaly detection:

- **LKE and ephemeral clusters** support AI-driven security workloads for large-scale scanning operations, inference pipelines, and threat modeling simulations.
- **Akamai's edge network** provides geographically distributed testing capabilities for latency-sensitive threat analysis and performance validation.
- **Compute, storage, and networking components** form a tightly coupled infrastructure supporting high-throughput data flows and automated resource management, enabling real-time intelligence delivery at global scale.

This setup has produced measurable benefits: faster and more consistent threat intelligence, superior operational reliability, and improved detection accuracy. Customers report quicker updates and enhanced responsiveness, validating the impact of ephemeral, Kubernetes-based orchestration in conjunction with a globally distributed edge network.

Delivering multi-region security intelligence in minutes

One illustrative example involved a multi-region digital footprint assessment for a client. spiderSilk's platform scanned thousands of domains simultaneously using ephemeral clusters across Akamai's edge network. The geographically distributed architecture enabled accurate, low-latency intelligence gathering from multiple vantage points.

"We were able to simulate realistic threat scenarios with Akamai — something that would have been impossible with a single-region setup. Plus, we realized other huge benefits," said Shah.

LKE's ephemeral clusters reduced processing times from multiple hours to under 30 minutes for global threat assessments, and decreased infrastructure costs by 30%–40%. They also enabled spiderSilk to reclaim 20%–30% of engineering time each week for higher-value activities such as model development and threat intelligence.

The combination of scalable infrastructure, edge native capabilities, and responsive support makes Akamai a true strategic partner in our mission to redefine cyber-AI.

Furqan Shah
Senior DevOps Lead, spiderSilk

High-touch support inspires operational confidence

Akamai's **Enhanced Compute Support** — offering expert assistance, faster response times, and a centralized support team — has been critical in sustaining spiderSilk's operational confidence. Direct access to engineers allows rapid troubleshooting, optimized scaling, and proactive guidance for high-stakes workloads.

For instance, during the multi-region assessment, Akamai helped the spiderSilk team adjust cluster configurations in real time, avoiding potential service disruption and ensuring on-schedule completion. "This level of support differentiates Akamai from hyperscaler alternatives, where ticket-based workflows often slow response times," explained Shah.

For spiderSilk, the combination of predictable infrastructure, global reach, and expert support enables reliable operations, faster deployment, and confidence during periods of rapid growth or complex AI processing.

Preparing for the next stage of growth

Following CPX Holding's acquisition of spiderSilk, the company has continued expanding its platform. Processing capacity has grown 25%–30% to handle more complex AI and scanning workloads, additional regions have been added to reduce latency and improve coverage, and the customer base has grown nearly 40%.

Looking ahead, spiderSilk plans to integrate GPU-accelerated workloads for deep learning, real-time inference, graph analytics, and advanced threat correlation. "Expanded edge coverage and enhanced orchestration tools will further improve global performance, maintain cost efficiency, and enable continuous experimentation for our AI-driven cybersecurity innovation," said Shah.

By scaling in step with spiderSilk's evolving needs, Akamai Cloud provides the foundation for real-time intelligence, operational resilience, and global cybersecurity leadership. "As our workloads grow and diversify, Akamai allows us to focus entirely on innovation and delivering actionable intelligence for our customers worldwide," concluded Shah.

>spiderSilk;

spiderSilk, founded in 2019 by Rami El Malak and Mossab Hussein, is the region's pioneering cyber-AI product company, renowned globally for its innovative AI-powered cybersecurity solutions. spiderSilk provides advanced defensive technologies, including its flagship exposure management platform, Resonance, autonomous SOC agents, and an internet-wide intelligence platform. Headquartered in the UAE, spiderSilk serves a global customer base across the UAE, Saudi Arabia, and the United States, with a dedicated R&D center continuously driving innovation. For more information, visit spidersilk.com.

© 2026 Akamai Technologies | Support |
| Published 03/26