splunk> a CISCO company

CUSTOMER STORY

Mitiga Reduces Data Collection Failure Detections from a Day to Hours

Key Challenges

Mitiga's customers count on them to catch and stop cloud attacks. That only works if the logs powering the detection of clousd attacks and other threats continue to flow. When collections experienced a volume drop, it signaled a blind spot at the worst possible moment. While these drops were noticed, prior to Cisco confirming and validating a drop could take up to 24 hours.

Key Results

Mitiga used the Cisco Deep Time Series Model (CDTSM) to power anomaly detection on its collection monitoring, helping the team catch volume drops sooner and confirm issues faster. CDTSM, a zero-shot foundation model for time series prediction, forecasts the expected volume for each collection job, flagging any volume drop outside the expected range. Running in the Splunk AI Toolkit across 10,000+ collection jobs, CDTSM reduces the time detecting broken collections from a full day down to a few hours.

mitiga

Industry: Security (Cloud and SaaS Detection and Response)

Solutions: Mitiga - Agentic Runtime Security across cloud, SaaS, identity, and AI

Products : Splunk AI Toolkit, Cisco Deep Time Series Model, Splunk MCP Server

Outcomes

- A few hours to detect a collection issue, down from ~24 hour
- ~95% reduction in false positive alerts
- ~4x improvement in mean time to detect (MTTD) for collection issues

An early adopter operating at scale

Founded in 2019, Mitiga is an Agentic Runtime Security platform providing comprehensive protection across cloud, SaaS, identity, third-party services, and AI. As an early adopter of the Cisco Deep Time Series Model, Mitiga uses the technology to power the anomaly-detection capabilities that form the backbone of its security monitoring, serving major enterprises globally including Blackstone, Redis, and New American Funding.

Initially launched as a cloud incident-response company, its founders saw that prevention alone wasn't enough, and that legacy tools weren't built to investigate and stop active cloud attacks. Migita closed that gap with Helios AIDR, its AI-native detection and response engine, pairing AI-driven detection and automation with a Cloud Security Data Lake holding more than 1,000 days of activity. Mitiga calls the outcome "Zero-Impact Breach Prevention," which contains attacks before they reach the business.

Delivering on that mission depends on reliably collecting logs from customers' cloud environments around the clock. Mitiga streams those logs 24/7 through dedicated pipelines ("jobs"). With more than 10,000 collection jobs running continuously, the health of that data pipeline is critical to ensuring that coverage remains intact. If a collection fails, the data stops flowing when customers can least afford it.

"Identifying a collection issue through volume drops could take us up to 24 hours," said Amit Matalon, DevOps Engineer, Mitiga. "With anomaly detection from Cisco Deep Time Series Model, we now catch those volume drops within hours. Reliable data collection is a critical part of the platform, so closing that gap mattered."

Catching volume drops early

For Mitiga, volume is one of the most critical signals. When event volume drops unexpectedly, it often indicates an upstream break. Previously, surfacing that break would take around 24 hours, far too long considering that data collection is a core part of Mitiga's platform.

The goal was to detect collection issues within a few hours, treating volume as a heartbeat signal, and doing it at scale across high-cardinality jobs. These challenges required anomaly detection with seasonality and baselining, smart alerting so that one outage produces one alert rather than a storm, and retention spanning high-resolution recent data, plus long-term aggregates.

Putting CDTSM to work on collection monitoring

As an early adopter of the Cisco Deep Time Series Model, Mitiga applied it to a problem that mattered deeply to its own platform: catching downward volume drops in log collection across tens of thousands of jobs. The model was a strong fit because, unlike traditional statistical approaches such as ARIMA or Prophet, it requires no custom training or fine-tuning. As a zero-shot foundation model, it generalizes across different data patterns out of the box, which suited Mitiga's complex, high-cardinality environment.

> **The model is pretrained to recognize what normal volume looks like. Instead of chasing every collector error, we only look at the drops it flags as outside the expected range. That is the difference between noise and something worth acting on.**
>
> Amit Matalon, DevOps Engineer, Mitiga

The detection runs as a single SPL search where each job's volume series is piped into the model with the AI Toolkit's apply command. The model returns a forecast of the expected range for that job and flags any observable volume drop outside of it. Because Splunk natively hosts CDTSM inside its Splunk platform, it processes that metric data without moving it to an external model or provider.

Mitiga tuned CDTSM to its own use case, configuring the detection sensitivity before it confirms an issue — for example, expecting a drop to persist across a few hours rather than reacting to a brief dip. With the model tuned to these specifications, quiet periods no longer look like outages, and only meaningful anomalies trigger an alert.

From experiment to production

What began as experimentation is now a steady, operationalized pipeline running continuously in production, watching that monitors the collection pipeline around the clock so the team can focus elsewhere.

"The Cisco Deep Time Series Model is pretrained to recognize what normal volume looks like. So instead of chasing every collector error, we only look at the drops it flags as outside the expected range," said Matalon. "That is the difference between noise and something worth acting on."

Mitiga's successful deployment demonstrates how advanced AI allows security teams to shift focus from chasing noise to acting on verified threats. And the company is continuing to refine the approach, exploring how to more deeply integrate AI/ML across its monitoring architecture.

As the company continues to scale, CDTSM-powered anomaly detection has become a core component of its infrastructure, proving that even at high volumes, machine-speed monitoring provides a strong competitive advantage.