DeepTempo

Case Study

LogLM-Powered Cybersecurity: Major Telco Implementation

How DeepTempo Performed real-time attack-intent detection

Executive Summary

A major telecommunications company evaluated DeepTempo's LogLM-powered solution to enhance customer cybersecurity. The initiative aimed to provide proactive threat detection for over 70 million customers, alerting them to potential device attacks. The results led to a paid production pilot at the Telco.

Business Challenges

- Avoid alert fatigue with low-noise detection
- Detect attacks with minimal false positives and false negatives
- Identify both known and novel behavior in near real time
- Operate fully on cloud-native infrastructure without heavy CPU/GPU use
- Require no constant human tuning
- Adapt instantly as traffic patterns shift
- Strengthen detection across the footprint with DeepTempo’s LogLM foundation model, built for this scale and complexity

Implementation & Testing

To understand how DeepTempo’s LogLM performs in real carrier environments, the operator executed a structured three-phase evaluation that tested accuracy, scalability, and generalization under real-world network conditions.

Solution Approach

Phase 1: Fine-Tuning on a Small Sample

In the first phase of the evaluation, the operator ran DeepTempo's LogLM with light fine-tuning on a small slice of their traffic distribution. This produced strong results across both benign and malicious activity. The model delivered an F1 score of 99, reflecting how well a fine-tuned model can adapt to a known environment - though we note that this level of precision exceeds what is typically expected in full production at national scale.

Across most attack classes, LogLM delivered excellent mean time to detect (MTTD) and high detect rates. Two patterns, however, highlighted important edge cases for flow-only detection. Brute-force activity was consistently identified, but only after more than 100 seconds of observation, resulting in a slower MTTD than the operator's target. Short-sequence infiltration attacks (Infil) were the only category the model struggled to classify accurately. These events generate minimal observable signal at the flow level - a known limitation of this generation of the model. Our team is already working on architectural improvements to address sparse-pattern scenarios.

High-signal behaviors such as DDoS, injection, XSS, password abuse, scanning, and reconnaissance consistently produced enough flow-level signal for immediate detection. Two predictable limits emerged: brute-force attacks take longer to reveal intent, and very short infiltration attempts generate too little signal for reliable classification.
MTTD reflects when DeepTempo receives the NetFlow telemetry, so timing varies by environment.

Phase 2: Full-Scale Testing on Remaining Traffic

In the second phase, the telco applied the same model fine-tuned in Phase 1 to the remaining 96% of production traffic - without any additional tuning. This tested whether LogLM could maintain accuracy through true zero-shot generalization across massive, diverse, real-world datasets. LogLM delivered a high F1 score of 0.96, supported by a 0.99 false-positive rate and 0.93 false-negative rate, demonstrating strong precision even as traffic volume and distribution expanded.

Solution Approach

Across most attack types, both MTTD and detect rates remained consistently strong. Two known edge cases behaved as expected:

- Brute Force: Detection remained accurate but required more than 100 seconds of sequence data before crossing the confidence threshold, resulting in longer MTTD than targeted.
- Infiltration: This attack type again showed low detect rates and slower MTTD, reflecting the inherent challenge of extremely short, low-signal sequences in flow-only analysis.

This phase confirmed that LogLM's foundation-model architecture can scale effectively across national traffic volumes while retaining precision - even when operating entirely zero-shot.

Phase 3: Zero-Shot Generalization on New Datasets

In the final phase, LogLM was tested on two entirely new datasets - with no additional fine-tuning. These datasets differed meaningfully from Phase 1 and 2, containing fewer distinct attack types occurring at higher frequency, providing a strong test of LogLM's ability to generalize beyond the original distribution. Across nearly all attack classes, MTTD and detect rates remained strong. One expected edge case persisted: **Infiltration**

This attack type again produced a very low detect rate, though in this dataset LogLM identified the limited signal rapidly, resulting in a fast MTTD despite overall low detectability. This third test validated that the fine-tuned model from Phase 1 can sustain high performance across fully new, real-world traffic distributions without retraining - an essential requirement for carrier-scale deployment.

DeepTempo

Case Study

Results:

Consolidated Results

Across all three phases, LogLM demonstrated a clear and repeatable pattern of performance. High-signal attack types - DDoS, injection, XSS, password abuse, scanning, and reconnaissance - were consistently detected with strong accuracy and rapid MTTD, even as dataset size, traffic distribution, and environmental conditions varied. Two predictable limitations appeared in every phase:

- Brute-force attacks were identified reliably but not immediately, requiring additional sequences before the model's confidence threshold was met.
- Infiltration and sparse infill patterns remained the weakest class, not due to model failure but because these attacks produce almost no discernible NetFlow signal. DeepTempo's research team is already developing targeted approaches to address this known challenge.

The most important finding: the model was not "tuned into" a single dataset — it retained high precision and low error rates even when applied to new traffic it had never seen before.

Test Phase	F1 Score	False Positives	False Negatives
Fine-tuned evaluation Dataset 1 (Phase 1 & 2)	98%	1.3%	0.5%
Zero-shot evaluation Dataset 1 (Phase 3)	90%	0.9%	5.9%
Zero-shot evaluation Dataset 2 (Phase 3)	73%	1.6%	2.2%

Business Outcomes

- Consistenly low MTTD, often identifying a concering pattern within the first sequence.
- Most attack classes were detected instantly, at high precision, with low false positives.
- Extremely high 90%+ zero shot performance in demanding environment.
- 96–98% F1 on new datasets, validating its ability to maintain precision when applied to unseen traffic.
- Validated feasibility of agent-free, large-scale attack detection.
- Demonstrated strong performance on typical attack patterns.
- Showed above expectation zero shot performance on a similar, but different data set.

Next Steps

Many Telcos, Service Providers, and Enterprises are challenged to identify attacks and are unable to use signature based detections. Many of these organizations also possess flow logs - often they are required to retain them by regulators and for use in forensics circumstances. The DeepTempo LogLM - and related software - is able to both see attacks rapidly and able to adapt to different environments, a capability that has been augmented through the use of patent pending Active Learning.

Please contact DeepTempo to evaluate the LogLM based solution; it can deploy in minutes on standard infrastructure and the DeepTempo team can assist in your evaluation, adaptation, and deployment. You can also utilize the LogLM as a part of the DeepTempo NativeApp on Snowflake. In addition to improving the protection of your enterprise, you can save funds by reducing the direct and indirect costs of false positives.

About DeepTempo

We're building the collective defense platform necessary to defend our businesses and critical infrastructure from the new breed of adversaries armed with the power of AI and the cloud. Together, we shall reclaim the advantage.

DeepTempo: The AI for Threat Detection.

Learn more: https://www.deeptempo.ai/

Follow us on LinkedIn: https://www.linkedin.com/company/deeptempo/