Simply. U+

ITU-T X.509 DAY · SESSION 3 · GENEVA · 12 MAY 2026

AI/ML-Proofing the SOC

A CSIRT Playbook for Threat Hunting and Incident Response in Mobile Networks

KIM YOUNG WOOK

Cybersecurity Incident Response Team · LG U+

The SOC is the first place AI/ML threats become real — and the last place that gets **the playbook update.**

How AI changes the adversary

Three things, simultaneously. The combination is what changes everything.

TIME

27 seconds

Fastest eCrime breakout time. 2026 CrowdStrike Global Threat Report. The defender's clock and the attacker's clock are no longer measured in the same units.

SCALE

Thousands in parallel

A single human operator runs 2–3 social engineering targets at once. AI-enabled adversaries surged 89% year-over-year — automation, not headcount, drove the increase. (2026 CrowdStrike GTR)

REALISM

Linguistic anomaly: gone

Generative models defeat the heuristics human analysts used for two decades. 'Check for grammar errors' is now actively misleading advice.

Four patterns the SOC already sees

Not predictions. Not red team exercises. Live tickets in our queue.

01

AI-Generated Social Engineering

Voice, email, SMS — linguistic detection signal is gone. Hunt has moved to behavioral signals.

02

Reconnaissance at Machine Speed

Scan-and-pair attacks compress disclosure-to-exploitation from weeks to hours.

03

Polymorphic Phishing Infrastructure

Domains and pages regenerate daily. The IOC-sharing foundation of CTI is eroding.

04

Adaptive Adversarial Behavior

Attacks observe defender response and adapt. SIEM rules now have shorter half-lives.

The CSIRT playbook, updated

Six stages. Each one rewritten for an adversary that does not work in human time.

STAGE	TRADITIONAL	UPDATED →
Detection	Signature-based + anomaly enrichment	**Behavioral baselines first, signatures second**
Triage	Volume the team can read	**AI-assisted prioritization; humans decide**
Containment	Human-time response (minutes)	**Pre-authorized automated containment**
Eradication	Remove known IOC artifacts	**Validate that behavioral pattern stopped**
Recovery	Validate the system is clean	**Brief humans: expect impersonation next**
Lessons	Document what happened	**Document near-misses and trajectories**

Threat hunting in the AI era

Hunting becomes more important, not less. It does not depend on yesterday's signatures.

HYPOTHESIS

Capability-anchored hunting

We construct hunt queries for attacker capabilities we believe are now possible — and run them before evidence appears in the wild.

TIME HORIZON

From batch to streaming

Hunts run continuously over short windows to catch the rapid adaptation cycle of AI-augmented attacks.

CADENCE

Hunt → Rule in 48 hours

Successful hunts produce detection rules within two days, knowing those rules will themselves have short half-lives.

> ***We are hunting for the attack we have not yet observed.***

Hunt cadence, hypothesis throughput, time-to-rule — none of these metrics existed in our SOC three years ago.

The CTI function, re-imagined

Three intelligence sources we now track that we did not three years ago.

01

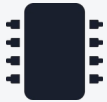

AI capability progression

What can current open-weight models actually do offensively? When a model release lowers the bar for an attack class, that is intelligence with operational consequence.

02

Underground forums & Telegram

Threat actor recruitment, tooling marketplaces, and AI-augmented attack discussions on Telegram, dark web forums, and access broker channels. The chatter often precedes campaigns by weeks.

03

Behavioral indicators

IPs, domains, hashes — atomic indicators decay fast. Behavioral patterns, sequencing, target selection logic — these are more durable against adaptation.

The next incident will arrive at human speed —

until it doesn't.

Thank you.