Cyber Threat Detection & Response

AI/ML-Based Threat Intelligence and Anomaly Detection Systems

1. Overview

Traditional signature-based security (looking for known "fingerprints" of malware) is no longer enough to stop modern cyber-attacks. Today's threats are polymorphic, fileless, and often powered by AI itself. This challenge tasks you with building an "Intelligent Immune System" for digital infrastructure—a system that uses Machine Learning to spot "weird" behavior before it turns into a full-scale breach.

2. The Challenge

Participants must develop an automated system capable of analyzing massive streams of data (network logs, user behavior, or system calls) to identify and neutralize threats in real-time. The goal is to reduce the "Mean Time to Detect" (MTTD) and "Mean Time to Respond" (MTTR) using intelligent automation.

3. Key Themes & Areas of Focus

To narrow down their solutions, encourage participants to look into:

- **User and Entity Behavior Analytics (UEBA):** Identifying compromised accounts by spotting deviations from a user's "normal" activity patterns.
- **Predictive Threat Intelligence:** Scraping the dark web or open-source intelligence (OSINT) to predict incoming attack vectors before they hit.
- **Automated Incident Response (SOAR):** Designing "Playbooks" where the AI doesn't just find the threat but automatically isolates the affected server or resets credentials.
- **Explainable AI (XAI) in Security:** Ensuring the system doesn't just say "Threat Detected," but explains *why* it flagged the activity so human analysts can trust the decision.

4. Expected Deliverables

- **Detection Engine:** A functional ML model (trained on datasets like CIC-IDS or similar) that can classify traffic as benign or malicious.
- **Real-time Dashboard:** A visual interface that displays live alerts, threat severity levels, and origin points.

- **Response Script/Action:** A demonstration of the system taking a "counter-measure" (e.g., automatically updating a firewall rule or killing a malicious process).
- **Technical Documentation:** A breakdown of the features used for training and how the model handles "False Positives."

5.Evaluation Criteria

- **Detection Accuracy (Precision & Recall):**We are looking for a balance between catching every threat and avoiding "The Boy Who Cried Wolf." How well does your model minimize False Positives while ensuring no critical "Zero-Day" attacks slip through?
- **Speed of Response (Latency)**:In cybersecurity, seconds matter. Your system will be evaluated on how quickly it processes incoming data and how fast it can trigger an automated response without human intervention.
- **Robustness against Adversarial AI:**Smart hackers try to "poison" the data to trick the AI. We want to see if your system is resilient—can it still detect a threat even if the attacker is trying to blend in or hide their patterns?
- **Adaptability & Learning:**A static model is a dead model. We are looking for systems that can learn from new data over time, adapting to shifting attack patterns without requiring a complete manual retraining of the engine.