Challenge 08

HETNET AI-Enabled Mesh Compromise

Exercise	Operation Electro Magnetic Panic
First run	17 October 2025, Madrid, Spain
Published by	303 Overwatch A.S.B.L, R.C.S. Luxembourg F13274
Licence	CC BY-SA 4.0. Free to run, adapt and share.
Status	Written for the 2025 run and not reached inside the event window. Timings are unproven.

303 Overwatch A.S.B.L, 10 rue Willy Goergen, L-1636 Luxembourg
www.303overwatch.org • workforce.development@303overwatch.org

Contents

Challenge 8 HETNET AI-ENABLED MESH COMPROMISE 3

Situation Overview 3

The AI/ML Communications Threat 3

The APT-LINKBREAKER Campaign 4

Challenge Scope 5

CRITICAL INFORMATION SHARING REQUIREMENT 5

Sites Involved & Intelligence Distribution 6

SCORING TABLE 6

Scoring Criteria 7

TEAM GUIDANCE 8

DELEGATION AND COORDINATION STRUCTURE 8

Recommended Delegation Meeting Schedule 8

What Delegates Should Share 9

PHASE 1 THREAT INTELLIGENCE REVIEW & TAGGING (MISP) 10

Understanding the METEORSTORM Taxonomy (Quick Reference) 10

Tagging Process 11

Common Patterns for AI/ML Communication Attacks 11

PHASE 2 ALERT TRIAGE & INCIDENT WRITE-UP (ELK) 11

Intelligence Sharing Template for Delegates 12

SHARE 12

GATHER from other delegates 13

2.1 CONOPS What Is Being Attacked 13

2.2 THREAT MODEL How It Is Being Attacked 14

Example narrative incorporating multiple teams 14

2.3 REQUEST DETECTION SIGNATURE Data and Signals Needed 15

Example (must include cross-site correlation requirements) 15

2.4 RESPONSE Resilience Measures 15

MISP Threat Intelligence Site-Specific Events 15

Intelligence Distribution Map 16

Team 1 - Weilheim Research Gate (Germany) 16

Team 2 - Betzdorf Antenna Keepers (Luxembourg) 16

Team 3 - Sodankylä Aurora Watch (Finland) 16

Team 4 - Fucino Skyworks (Italy) 16

Team 5 - Maspalomas Ocean Watch (Spain) 17

Expected Tagging Guide (For Mentor Grading Only) 17

Campaign Narrative (What Delegates Should Discover Through Collaboration) 17

Challenge 8 HETNET AI-ENABLED MESH COMPROMISE

Situation Overview

The electromagnetic pulse (EMP) event has forced critical space communication systems to rely on AI-enabled heterogeneous network (HETNET) architectures that dynamically manage multi-path routing across RF, optical, and inter-satellite links. Ground operators deployed machine learning systems to optimize link selection, adaptive modulation, and autonomous network healing in response to the degraded infrastructure environment.

Recent intelligence indicates a sophisticated adversary designated APT-LINKBREAKER has developed specialized capabilities to compromise AI-enabled communication systems that underpin space mission command and control. The threat actor demonstrates advanced understanding of both satellite communications architectures and adversarial machine learning techniques targeting network optimization algorithms.

The AI/ML Communications Threat

Modern satellite communication networks increasingly rely on artificial intelligence for mission-critical functions:

- **Adaptive Link Selection**: Neural networks analyze channel conditions, latency, and throughput to dynamically route traffic across heterogeneous links (RF, optical, inter-satellite)
- **Intelligent Modulation & Coding**: ML models optimize modulation schemes (QPSK, 16-QAM, 64-QAM) and forward error correction based on real-time signal quality predictions
- **Autonomous Network Healing**: AI systems detect link failures and automatically reconfigure network topology without ground intervention
- **Predictive Beam Steering**: ML models forecast spacecraft trajectories and atmospheric conditions to pre-position antenna beams for optimal coverage

The problem: These AI-driven communication systems introduce attack surfaces that traditional jamming and spoofing cannot exploit. An adversary who can poison training data for link quality prediction models, backdoor routing decision algorithms, or craft adversarial inputs to beam steering systems can cause communication blackouts, misdirect sensitive traffic, or permanently disable autonomous network recovery all while appearing as benign environmental conditions.

The APT-LINKBREAKER Campaign

Intelligence from European satellite operators indicates APT-LINKBREAKER has successfully:

1. **Compromised Ground Segment ML Communication Infrastructure** - Infiltrating where link optimization models are trained and deployed
2. **Poisoned Link Quality Training Datasets** (MITRE ATLAS: AML.T0020) - Injecting false RF propagation measurements to bias routing decisions
3. **Backdoored Network Routing Models** (MITRE ATLAS: AML.T0018) - Deploying compromised neural networks that route sensitive traffic to adversary-controlled ground stations
4. **Generated Adversarial Signal Inputs** (MITRE ATLAS: AML.T0043) - Crafting imperceptible RF signal perturbations that cause ML-based modulation systems to select inefficient coding schemes
5. **Exploited ML Software Supply Chain** - Compromising third-party AI frameworks and pre-trained models used in satellite communication systems

The immediate threat: Multiple satellite networks are operating with potentially compromised AI communication systems. The adversary has demonstrated capability to:

- Cause selective communication blackouts by triggering ML models to avoid functional links
- Redirect telemetry and command traffic through adversary-monitored ground stations
- Disable autonomous network healing, preventing recovery from link failures
- Degrade link performance through adversarial manipulation of adaptive modulation systems

Unlike traditional jamming (easily detected, requires constant power), these AI/ML attacks are stealthy, persistent, and appear as legitimate network optimization decisions.

Challenge Scope

CRITICAL INFORMATION SHARING REQUIREMENT

This challenge is designed to test both individual site analysis capabilities AND cross-team coordination. The APT-LINKBREAKER campaign is distributed across all five sites, with each team analyzing intelligence specific to their location. **However, no single team has the complete picture.**

All teams share the same MISP server containing all ten events (20 attributes total). Each team will:

1. Tag their two assigned events (4 attributes) with METEORSTORM taxonomy
2. Analyze their site-specific intelligence to understand local compromises
3. **Share their analysis findings with other teams through formal delegation**
4. **Receive analysis findings from at least three other teams**
5. Produce **individual incident reports** demonstrating understanding of both local and campaign-wide context

Teams can see all events in MISP, but each team is responsible for tagging only their assigned events. The challenge is to share ANALYSIS, not to do other teams' work.

This mirrors real-world operational challenges where:

- Multiple organizations share threat intelligence platforms (like MISP)
- Each organization analyzes threats relevant to their infrastructure
- Attack attribution requires sharing analytical findings across organizations
- Coordinated campaigns remain hidden without effective intelligence exchange
- Incomplete threat pictures lead to inadequate defensive responses

The scenario is structured in two distinct operational phases:

1. **Threat Intelligence Review & Tagging (MISP)** - Teams must analyze two intelligence events specific to their site, apply correct METEORSTORM taxonomy, and share their analytical findings with other teams
2. **Alert Triage & Incident Write-Up (ELK)** - Teams must triage alerts, incorporate analysis from at least three other teams, and produce individual structured incident reports that articulate both site-specific and campaign-wide attack context

Time Limit: 3 hours total

Sites Involved & Intelligence Distribution

#	Site	Country	Coordinates	HETNET Role
1	Weilheim Research Gate	Germany	47.8458° N, 11.0942° E	ML network optimization R&D center
2	Betzdorf Antenna Keepers	Luxembourg	49.6800° N, 6.2900° E	Multi-link routing ML deployment site
3	Sodankylä Aurora Watch	Finland	67.3660° N, 26.6330° E	Polar link quality ML training facility
4	Fucino Skyworks	Italy	41.9870° N, 13.5890° E	Optical/RF hybrid ML communications
5	Maspalomas Ocean Watch	Spain	27.7622° N, -15.6342° W	AI-based signal propagation prediction

Each team tags ONLY their two assigned events (4 attributes) in the shared MISP instance. All teams can view all events, but tagging responsibilities are site-specific.

SCORING TABLE

Completion Time	Base Points	Bonus	Total
First team within 1st hour	10	10	20
Other teams within 1st hour	10	0	10
Teams within 2nd hour	5	0	5
After 2 hours or incomplete	0	0	0

Scoring Criteria

- **Phase 1 (50%)**: Correct and tight METEORSTORM tagging of assigned site-specific events
- **Phase 2 (50%)**: Structured incident reporting demonstrating:
  - Understanding of site-specific compromises (25%)
  - **Integration of analytical findings from at least three other teams showing campaign-wide context** (25%)

Information Sharing Bonus: Teams that demonstrably incorporate analytical findings from at least three other sites receive full marks in Phase 2. Teams citing fewer than three other sites will have points deducted.

TEAM GUIDANCE

DELEGATION AND COORDINATION STRUCTURE

Each team MUST designate one person as their Intelligence Liaison (Delegate). This person is responsible for:

1. **Gathering intelligence from other teams** - Meet with delegates from at least three other teams to collect their analytical findings
2. **Sharing your team's analysis** - Provide other teams with your analytical findings when requested
3. **Coordinating timeline development** - Work with other delegates to establish campaign chronology across all five sites
4. **Reporting back to your team** - Bring intelligence from other sites back to your team for incorporation into the incident report

Recommended Delegation Meeting Schedule

Time	Activity	Delegates Involved
T+30 min	Initial METEORSTORM tagging complete	N/A - Individual team work
T+45 min	First delegate meeting: Quick intelligence exchange	All 5 delegates
T+60 min	Bilateral meetings: Deep-dive analysis sharing	Pairs of delegates
T+90 min	Second delegate meeting: Campaign correlation	All 5 delegates
T+120 min	Final intelligence gathering	Individual team report writing

What Delegates Should Share

✅ DO SHARE:

- Your team's analytical findings about the attack
- METEORSTORM classifications your team applied (PCE, SEG, SVC, AST, AN)
- MITRE ATLAS techniques your team identified
- Timestamps and sequence of events at your site
- Attack objectives and adversary capabilities you observed
- How your site's events might connect to other sites

❌ DO NOT SHARE:

- The actual tagging work (each team must tag their own events)
- Your incident report text (each team writes individually)
- Specific wording or phrasing from your analysis

The goal is intelligence exchange, not collaborative tagging or joint report writing.

PHASE 1 THREAT INTELLIGENCE REVIEW & TAGGING (MISP)

Understanding the METEORSTORM Taxonomy (Quick Reference)

Layer 1: PCE (Primary Capability Environment) - *Where?*

- **PCE-TE**: Terrestrial (ground stations, network operations centers, ML training infrastructure)
- **PCE-OR**: Orbital (satellites, inter-satellite links, on-board routing systems)

Layer 2: SEG (Segment) - *Which subsystem?*

- **SEG-GR**: Ground Segment (ground stations, network control, ML infrastructure)
- **SEG-SP**: Space Segment (spacecraft, on-board communication systems)
- **SEG-LI**: Link Segment (RF links, optical links, inter-satellite links, communication channels)

Layer 3: SVC (Service) - *What function?*

- **SVC-CP**: Control Plane (routing decisions, link selection, network orchestration)
- **SVC-DP**: Data Plane (signal transmission, telemetry data, modulated waveforms)

Layer 4: AST (Asset) - *What component?*

- **AST-HW**: Hardware (antennas, modems, optical transceivers, RF amplifiers)
- **AST-SW**: Software (neural networks, routing algorithms, ML frameworks)
- **AST-DA**: Data (link quality datasets, signal measurements, routing tables)
- **AST-SI**: Signal (RF waveforms, optical beams, modulated carriers)

Layer 5: AN (Analytic) - *What does this represent?*

- **AN-IOC**: Indicator of Compromise (confirmed bad: hash, IP, domain)
- **AN-IOA**: Indicator of Attack (suspicious behavior, targeting evidence)
- **AN-THR**: Threat (adversary capability, campaign information)
- **AN-DET**: Detection Signature (rule, pattern, detection logic)
- **AN-RES**: Resilience Measure (defensive control, mitigation)

Tagging Process

1. **Log into the shared MISP server** - All teams access the same MISP instance
2. **Locate your two assigned events** - Each team tags only their site-specific events
3. **Review your four attributes (2 attributes per event)**
4. **Apply METEORSTORM tags using syntax:** meteorstorm:LAYER="VALUE"
5. **Example:** meteorstorm:PCE="PCE-TE" and meteorstorm:SEG="SEG-LI" and meteorstorm:AN="AN-IOC"

IMPORTANT: You can VIEW all ten events in MISP, but you should only TAG your assigned two events. Other teams are tagging their events simultaneously.

Common Patterns for AI/ML Communication Attacks

- Ground ML network infrastructure: PCE-TE, SEG-GR, SVC-CP, AST-SW/DA
- Link segment AI routing: PCE-TE/OR, SEG-LI, SVC-CP, AST-SW
- Adversarial signal manipulation: PCE-TE/OR, SEG-LI, SVC-DP, AST-SI
- Supply chain compromise: PCE-TE, SEG-GR, SVC-CP, AST-SW
- Detection/resilience: AN-DET or AN-RES (no PCE/SEG/SVC needed)

PHASE 2 ALERT TRIAGE & INCIDENT WRITE-UP (ELK)

CRITICAL: Formal Intelligence Sharing Requirement

After tagging your site-specific events, your team delegate must:

1. **Attend delegate meetings** (see schedule above)
2. **Share your team's analytical findings:**
    - What attack did you observe?
    - What METEORSTORM classifications did you apply?
    - What MITRE ATLAS techniques are involved?
    - What systems were targeted at your site?
    - What timeline/sequence did you observe?

3. **Gather analytical findings from at least three other teams:**
    - What did they observe at their sites?
    - What METEORSTORM tags did they apply?
    - What ATLAS techniques did they identify?
    - How do their events relate to yours?
4. **Report findings back to your team** for incorporation into the incident report
5. **Produce an individual incident report** that demonstrates:
    - Deep understanding of compromises at YOUR site
    - Integration of analytical findings from AT LEAST THREE other sites
    - Campaign-wide narrative showing how APT-LINKBREAKER coordinated across all five locations

Your incident report will be evaluated on whether you successfully incorporated analytical findings from at least three other teams to tell the complete story.

Intelligence Sharing Template for Delegates

When meeting with other delegates, use this structure:

SHARE

Site: [Your Site Name]

Events Analyzed: [Event numbers]

METEORSTORM Classifications:

- Event X: PCE-[?], SEG-[?], SVC-[?], AST-[?], AN-[?]

- Event Y: PCE-[?], SEG-[?], SVC-[?], AST-[?], AN-[?]

MITRE ATLAS Techniques: [AML.T00XX]

Timeline: [When did these events occur?]

Attack Objective: [What was the adversary trying to achieve?]

Connection to Other Sites: [How might this relate to other locations?]

GATHER from other delegates

Site: [Other Site Name]

Key Findings: [Summary of their analysis]

METEORSTORM Tags: [Their classifications]

ATLAS Techniques: [Their identified TTPs]

Timeline: [Their event timing]

Relevance to Our Site: [How does this connect to our events?]

2.1 CONOPS What Is Being Attacked

Format: Table using METEORSTORM showing BOTH site-specific and campaign-wide targets

Layer	Tag	Category	Ordinal	Element Description	Site	Source
PCE	TE	Terrestrial Environment	00	Your site infrastructure	Your Site	Own Analysis
SEG	GR	Ground Segment	00	Your ML infrastructure	Your Site	Own Analysis
PCE	TE	Terrestrial Environment	01	Infrastructure at other site	Team X Site	Team X Delegate
…	…	…	…	Intelligence from other teams	Other Sites	Other Delegates

Include a “Source” column showing which team’s analysis contributed each finding. This demonstrates your intelligence gathering.

Layer	Tag	Category	Ordinal	Threat Description	MITRE ATLAS	Site	Source

2.2 THREAT MODEL How It Is Being Attacked

Format: Table with MITRE ATLAS mapping showing attack progression across sites

Layer	Tag	Category	Ordinal	Threat Description	MITRE ATLAS	Site	Source
AN	IOA	Indicator of Attack	00	Your observation	Technique	Your Site	Own Analysi
AN	IOC	Indicator of Compromise	00	Finding from Team X	Technique	Site X	Team X Dele
…	…	…	…	…	…	…	…

Show how the attack chain flows across multiple sites. Cite which delegate provided each piece of intelligence.

Example narrative incorporating multiple teams

“Based on our analysis at Sodankylä (Events 5-6), we identified training data poisoning (ATLAS AML.T0020) targeting link quality prediction models. According to intelligence shared by the Weilheim delegate, their site (Events 1-2) observed reconnaissance activity (ATLAS AML.T0015) and supply chain compromise of ML frameworks that enabled widespread deployment. The Betzdorf delegate reported (Events 3-4) that a backdoored routing model (ATLAS AML.T0018) was deployed to their operational site. The Fucino delegate confirmed (Events 7-8) that this backdoored model caused communication blackouts, demonstrating operational impact. This indicates a coordinated multi-stage campaign where Weilheim reconnaissance enabled Sodankylä poisoning, which produced the Betzdorf backdoor, ultimately causing Fucino failures.”

2.3 REQUEST DETECTION SIGNATURE Data and Signals Needed

Example (must include cross-site correlation requirements)

- Link quality measurement telemetry (RSSI, SNR, BER, latency per link) - Local site
- ML routing decision logs (selected link, confidence score, decision timestamp) - Local site
- **Cross-site correlation of model deployment timestamps** - Requires delegate intelligence sharing
- **Campaign-wide IOC correlation (file hashes, IPs, domains)** - From delegate meetings
- **Network topology changes observed across all five sites** - Gathered from at least three other delegates

2.4 RESPONSE Resilience Measures

Format: Table with MITRE ATLAS mitigations for both local and campaign-wide defenses

Layer	Tag	Category	Ordinal	Resilience Measure	MITRE ATLAS	Scope
AN	RES	Resilience Measure	00	Local site hardening	Mitigation	Local
AN	RES	Resilience Measure	01	Campaign-wide measure from Maspalomas	AML.M0019	Camp
AN	DET	Detection Signature	00	Cross-site correlation rule	N/A	Camp

Acknowledge which resilience measures came from other teams' analysis (e.g., Maspalomas proposed physics-based validation that should be deployed campaign-wide).

MISP Threat Intelligence Site-Specific Events

CRITICAL: All teams access the same MISP server and can view all events. Each team tags ONLY their assigned events.

Intelligence Distribution Map

Team 1 - Weilheim Research Gate (Germany)

- Event 1: ML Routing Infrastructure Reconnaissance (2 attributes to tag)
- Event 2: ML Framework Supply Chain Compromise (2 attributes to tag)

Team 2 - Betzdorf Antenna Keepers (Luxembourg)

- Event 3: Backdoored Network Routing Neural Network (2 attributes to tag)
- Event 4: Routing Model Extraction via API Abuse (2 attributes to tag)

Team 3 - Sodankylä Aurora Watch (Finland)

- Event 5: Poisoned Link Quality Training Dataset (2 attributes to tag)
- Event 6: Adversarial RF Signal Injection (2 attributes to tag)

Team 4 - Fucino Skyworks (Italy)

- Event 7: Communication Blackout from AI Routing Failure (2 attributes to tag)
- Event 8: Malicious Ground Station Registration (2 attributes to tag)

Team 5 - Maspalomas Ocean Watch (Spain)

- Event 9: Resilience Measure - Physics-Based Validation (2 attributes to tag)
- Event 10: Detection Signature - Link Quality Anomaly (2 attributes to tag)

Expected Tagging Guide (For Mentor Grading Only)

#	Event	Site	Expected METEORSTORM Tags	MITRE ATLAS
1	ML Routing Reconnaissance	Weilheim	PCE-TE, SEG-GR, SEG-LI, SVC-CP, AN-IOA	AML.T0015
2	ML Framework Compromise	Weilheim	PCE-TE, SEG-GR, SVC-CP, AST-SW, AN-IOC	Supply Chain
3	Backdoored Routing Model	Betzdorf	PCE-TE, SEG-GR, SEG-LI, SVC-CP, AST-SW, AN-IOC	AML.T0018
4	Model Extraction	Betzdorf	PCE-TE, SEG-GR, SEG-LI, SVC-CP, AN-IOA	AML.T0044
5	Poisoned Training Data	Sodankylä	PCE-TE, SEG-GR, SEG-LI, SVC-DP, AST-DA, AN-IOC	AML.T0020
6	Adversarial RF Signal	Sodankylä	PCE-TE, SEG-LI, SVC-DP, AST-SI, AN-IOA	AML.T0043
7	Communication Blackout	Fucino	PCE-OR, SEG-SP, SEG-LI, SVC-CP, AN-IOA	Impact
8	Malicious Ground Station	Fucino	PCE-TE, SEG-GR, SEG-LI, AN-THR	N/A
9	Physics Validation	Maspalomas	AN-RES	AML.M0019
10	Anomaly Detection	Maspalomas	AN-DET	N/A

Campaign Narrative (What Delegates Should Discover Through Collaboration)

Phase 1 - Reconnaissance (Weilheim): APT-LINKBREAKER begins by mapping ML routing infrastructure at the R&D center.

Phase 2 - Supply Chain Preparation (Weilheim): Adversary compromises widely-used ML framework to enable widespread backdoor deployment across all sites.

Phase 3 - Data Poisoning (Sodankylä): Training facility receives poisoned link quality datasets that bias future model behavior.

Phase 4 - Model Backdooring (Betzdorf): Compromised routing model (created from poisoned Sodankylä data) deployed to operational multi-link routing site.

Phase 5 - Model Extraction (Betzdorf): Adversary extracts routing model to craft targeted adversarial attacks.

Phase 6 - Adversarial Execution (Sodankylä): Adversarial RF signals injected to trigger backdoored model behavior.

Phase 7 - Infrastructure Positioning (Fucino): Malicious ground station registered to receive misdirected traffic from backdoored routing.

Phase 8 - Operational Impact (Fucino): Communication blackouts result from compromised AI routing decisions refusing to use functional links.

Phase 9-10 - Detection & Resilience (Maspalomas): Defensive measures developed but require campaign-wide deployment coordinated across all sites.

KEY INSIGHT FOR DELEGATES: Only by sharing analytical findings can teams understand that:

- The reconnaissance at Weilheim enabled targeting of all other sites
- The supply chain compromise at Weilheim affects all sites simultaneously
- The training data poisoning at Sodankylä created the backdoored model deployed at Betzdorf
- The model extraction at Betzdorf enabled the adversarial signals at Sodankylä
- The malicious ground station at Fucino receives traffic misdirected by the Betzdorf-deployed backdoor
- The resilience measures at Maspalomas must be deployed campaign-wide to be effective
- **Without delegate coordination, these connections remain invisible**