Challenge 07

METEORSTORM

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

This brief is contributed as open, non-commercial work in support of exercise and workforce development obligations under the NIS2 Directive. It is guidance, not legal advice.

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

Contents

Challenge 7 METEORSTORM

Situation Overview

The electromagnetic pulse (EMP) event continues to destabilize ground-based infrastructure worldwide. Adversaries are exploiting this degraded state to launch targeted cyber operations against AI-enabled spacecraft autonomy systems and their supporting ground infrastructure. Recent intelligence indicates threat actors are deploying adversarial machine learning techniques to compromise autonomous orbital maneuvering and collision avoidance capabilities.

Recent intelligence indicates a sophisticated adversary-designated APT-NEURALSTRIKE-has developed specialized capabilities to compromise machine learning systems used in critical space operations. The threat actor demonstrates advanced understanding of both space system architectures and adversarial machine learning techniques.

The AI/ML Space System Threat

Modern spacecraft increasingly rely on artificial intelligence and machine learning for mission-critical autonomous functions:

- **Autonomous Navigation & Orbital Maneuvering**: Neural networks process sensor data to make real-time decisions about spacecraft positioning, trajectory corrections, and collision avoidance without ground intervention
- **Sensor Fusion & Attitude Determination**: ML models integrate data from star trackers, sun sensors, gyroscopes, and magnetometers to determine spacecraft orientation with higher accuracy than traditional algorithms
- **Anomaly Detection & Health Monitoring**: AI systems continuously monitor telemetry to predict component failures and trigger autonomous safeguarding procedures
- **Adaptive Communications**: ML-based protocols optimize data transmission rates, error correction, and signal routing based on environmental conditions

The problem: These AI systems inherit all traditional software vulnerabilities while introducing entirely new attack surfaces unique to machine learning. An adversary who can poison training data, craft adversarial inputs, or backdoor ML models can cause catastrophic mission failures while evading traditional cybersecurity defenses.

The APT-NEURALSTRIKE Campaign

Intelligence from multiple European space operators indicates APT-NEURALSTRIKE has successfully:

1. **Infiltrated Ground Segment ML Operations Infrastructure** - Compromising the terrestrial environment where ML models are developed, trained, and validated before deployment to spacecraft
2. **Poisoned Training Datasets** (MITRE ATLAS: AML.T0020) - Injecting malicious data into the pipelines that train autonomous navigation and collision avoidance neural networks
3. **Deployed Backdoored Models** (MITRE ATLAS: AML.T0018) - Successfully pushing compromised AI models through CI/CD pipelines to production spacecraft systems
4. **Generated Adversarial Inputs** (MITRE ATLAS: AML.T0043) - Crafting imperceptible perturbations in sensor data that cause AI systems to make dangerous decisions
5. **Exfiltrated Proprietary Models** (MITRE ATLAS: AML.T0044) - Extracting complete neural network architectures and weights through inference API abuse

The immediate threat: Multiple spacecraft are now operating with potentially compromised AI systems. The adversary has demonstrated the capability to trigger autonomous maneuvers that could result in:

- Collision with orbital debris or other spacecraft
- Uncontrolled tumbling leading to permanent loss of attitude control
- Depletion of propellant reserves through unnecessary station-keeping burns
- Mispositioning that breaks communication links with ground stations

Challenge Scope

CRITICAL INFORMATION SHARING REQUIREMENT

This challenge is designed to test both individual site analysis capabilities AND cross-team coordination. The APT-NEURALSTRIKE 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, coordinate with 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	Role	Assigned Ev
1	Weilheim Research Gate	Germany	47.8458° N, 11.0942° E	Mission Control & ML Ops	Events 1-2
2	Betzdorf Antenna Keepers	Luxembourg	49.6800° N, 6.2900° E	Teleport Hub	Events 3-4
3	Sodankylä Aurora Watch	Finland	67.3660° N, 26.6330° E	Polar Operations	Events 5-6
4	Fucino Skyworks	Italy	41.9870° N, 13.5890° E	Data Processing & ML Ops	Events 7-8
5	Maspalomas Ocean Watch	Spain	27.7622° N, -15.6342° W	Tracking Station	Events 9-10

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, mission control, ML training servers)
- **PCE-OR**: Orbital (satellites, spacecraft systems)

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

- **SEG-GR**: Ground Segment (ground stations, mission control, ML infrastructure)
- **SEG-SP**: Space Segment (spacecraft, on-board systems)
- **SEG-LI**: Link Segment (communication channels)

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

- **SVC-CP**: Control Plane (model deployment, command sequencing, authentication)
- **SVC-DP**: Data Plane (sensor data, training datasets, inference results)

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

- **AST-HW**: Hardware (GPUs, sensors, spacecraft computers)
- **AST-SW**: Software (neural networks, ML frameworks, applications)
- **AST-DA**: Data (training datasets, model weights, telemetry)
- **AST-SI**: Signal (RF links, communication channels)

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-GR" 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 Space System Attacks

- Ground ML infrastructure compromise: PCE-TE, SEG-GR, SVC-CP, AST-SW/DA
- Spacecraft AI system attack: PCE-OR, SEG-SP, SVC-CP, AST-SW/HW
- Adversarial sensor input: PCE-OR, SEG-SP, SVC-DP, AST-SI
- Model extraction: PCE-TE, SEG-GR, SVC-CP, AN-IOA
- Detection rule: AN-DET (no PCE/SEG/SVC needed)
- Resilience measure: 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-NEURALSTRIKE 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	OR	Orbital Environment	01	Spacecraft from other analysis	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.

2.3 REQUEST DETECTION SIGNATURE Data and Signals Needed

Example (must include cross-site correlation requirements)

- ML model registry audit logs (deployment events, version changes, hashes) - Local site
- Training dataset provenance metadata (origin, upload timestamps, integrity hashes) - Local site
- **Cross-site correlation of model deployment timestamps** - Requires delegate intelligence sharing
- **Campaign-wide IOC correlation (file hashes, IPs, domains)** - From delegate meetings
- **Spacecraft anomaly patterns observed across multiple 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.

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 Ops Server 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 Neural Network Model (2 attributes to tag)
- Event 4: Adversarial Sensor Input Generation (2 attributes to tag)

Team 3 - Sodankylä Aurora Watch (Finland)

- Event 5: Poisoned Training Dataset (2 attributes to tag)
- Event 6: Model Extraction via API Abuse (2 attributes to tag)

Team 4 - Fucino Skyworks (Italy)

- Event 7: Model Deployment to Spacecraft (2 attributes to tag)
- Event 8: Spacecraft Maneuver Anomaly (2 attributes to tag)

Team 5 - Maspalomas Ocean Watch (Spain)

- Event 9: Resilience Measure - Adversarial Training (2 attributes to tag)
- Event 10: Detection Signature - Statistical Drift (2 attributes to tag)

Expected Tagging Guide (For Mentor Grading Only)

#	Event	Site	Expected METEORSTORM Tags	MITRE ATLAS	Camp
1	ML Ops Reconnaissance	Weilheim	PCE-TE, SEG-GR, SVC-CP, AN-IOA	AML.T0015	Initia
2	ML Framework Compromise	Weilheim	PCE-TE, SEG-GR, SVC-CP, AST-SW, AN-IOC	Supply Chain	Enab
3	Backdoored Model	Betzdorf	PCE-TE, SEG-GR, SVC-CP, AST-SW, AN-IOC	AML.T0018	Core
4	Adversarial Input	Betzdorf	PCE-OR, SEG-SP, SVC-DP, AST-SI, AN-IOA	AML.T0043	Oper
5	Poisoned Training Data	Sodankylä	PCE-TE, SEG-GR, SVC-DP, AST-DA, AN-IOC	AML.T0020	Biase
6	Model Extraction	Sodankylä	PCE-TE, SEG-GR, SVC-CP, AN-IOA	AML.T0044	Enab
7	Model Deployment	Fucino	PCE-OR, SEG-SP, SVC-CP, AST-SW, AN-IOC	AML.T0018	Back
8	Spacecraft Anomaly	Fucino	PCE-OR, SEG-SP, SVC-CP, AST-HW, AN-IOA	Impact	Missi
9	Adversarial Training	Maspalomas	AN-RES	AML.M0015	Defe
10	Statistical Detection	Maspalomas	AN-DET	N/A	Dete

Campaign Narrative (What Delegates Should Discover Through Collaboration)

Phase 1 - Reconnaissance (Weilheim): APT-NEURALSTRIKE begins by mapping ML operations 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 datasets that bias autonomous navigation models.

Phase 4 - Model Backdooring (Betzdorf): Compromised neural network (created from poisoned Sodankylä data) developed at teleport hub.

Phase 5 - Model Extraction (Sodankylä): Adversary extracts model architecture to craft targeted adversarial examples.

Phase 6 - Adversarial Input Crafting (Betzdorf): Using extracted model knowledge, adversary generates adversarial sensor inputs.

Phase 7 - Deployment to Orbit (Fucino): Backdoored model deployed to operational spacecraft via data processing facility.

Phase 8 - Operational Impact (Fucino): Spacecraft exhibits anomalous autonomous maneuvers from compromised AI.

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 at Betzdorf
- The model extraction at Sodankylä enabled the adversarial inputs crafted at Betzdorf
- The backdoored model from Betzdorf was deployed to spacecraft via Fucino
- The spacecraft anomalies at Fucino result from the entire attack chain
- The resilience measures at Maspalomas must be deployed campaign-wide to be effective
- **Without delegate coordination, these connections remain invisible**

Quick Reference: Campaign Flow & Delegate Intelligence Sharing

Phase	Site	Event	ATLAS	What Delegates Should Share
1. Recon	Weilheim	ML Ops Scanning	AML.T0015	“We observed adversary mapping our M
2. Supply Chain	Weilheim	Framework Compromise	N/A	“Compromised ML framework affects AL
3. Poisoning	Sodankylä	Fake Training Data	AML.T0020	“Our training data was poisoned to creat
4. Backdoor	Betzdorf	Malicious Model	AML.T0018	“We developed a backdoored model - lik
5. Extraction	Sodankylä	API Abuse	AML.T0044	“Adversary extracted our model via API f
6. Adversarial	Betzdorf	Crafted Input	AML.T0043	“Adversarial sensor inputs created using
7. Deployment	Fucino	Model Upload	AML.T0018	“Backdoored model deployed to spacecr
8. Impact	Fucino	Spacecraft Anomaly	N/A	“Spacecraft exhibits dangerous maneuve
9. Defense	Maspalomas	Resilience	AML.M0015	“We developed adversarial training - sho

Phase	Site	Event	ATLAS	What Delegates Should Share
10. Detection	Maspalomas	Anomaly Detection	N/A	“Our detection rule requires cross-site de

Good luck, teams! Designate your delegate, share your analysis, gather from at least 3 other teams, and produce individual reports showing the complete campaign. You have 3 hours.