FULL SPECTRUM SPACE CYBERSECURITY PROFESSIONAL

MODULE WORKBOOK · STUDENT EDITION

Module 02 Contextualized Threat Modeling

The method behind the Day-2 threat catalogue, and every entry it produced, each anchored to the one element it targets.

COPYRIGHT AND LEGAL NOTICES

PUBLISHER	D2 Team CORP, a 501(c)(3) nonprofit
PROGRAM	Full Spectrum Space Cybersecurity Professional
DOCUMENT	Module 02 Contextualized Threat Modeling, module workbook, student edition
EDITION	Student edition · revised 2026-09-24
CONTACT	www.d2teamcorp.org

STUDENT EDITION

The teaching chapters and the complete worked deliverable. Your instructor holds the instructor edition, which adds the session plan and the marking guidance.

SCOPE AND STANDING OF THIS WORK · READ BEFORE JUDGING THE ENUMERATION

What this is. Kestrel Orbital's first application of the METEORSTORM framework, produced by three departments who had never used it before, in five working days. It is a deliberately scoped first pass over the telecommand path: the command-and-control chain from the operator console to execution on orbit.

Why that scope, and not more. The mandates chose it. Executive Order 14144 puts the command and control of an operational platform first, and the NIS2 Directive places a duty of care and a reporting clock on the ground operator. The fleet is already flying, so the first pass ran on the part that would end the mission if an adversary took it. Enumerating everything else first would have been thorough and useless.

What it is not, stated plainly. It is not an exhaustive model of the platform and does not claim to be. Elements off the telecommand path are not enumerated. Some assets group several physical units where those units share an owner, a build and a patch path. Gaps in the ordinal sequence mark elements the register tracks outside this scope. Every one of those is a recorded decision taken at the enumeration meeting, not an oversight discovered later.

How it grows. Ordinals are issued once and never renumbered, so everything enumerated here keeps its identifier as the register expands. The same eight-step procedure runs on the next scope, and the one after that. A first pass that names its boundary can be extended. One that hides its boundary has to be rebuilt.

The standard we hold ourselves to. A scope choice is stated openly. A defect is fixed, never narrated. If you find an element misnamed, a parent wrong, or a count that does not add up, that is a defect and we want to hear about it. If you find something outside the telecommand path missing, that is the scope, it was chosen deliberately, and it is on the roadmap. The two are not the same criticism, and this work is built so you can tell them apart.

CONTENTS

PART I

Orientation

What the day produces, and what each department brings to the table before any of it can be written.

CH 1 THE DAY IN ONE PAGE

Day 2 opens in the Security Operations Center, and this time the department is presenting, not being sold. Yesterday Satellite Operations and Satellite Design & Engineering reconciled their views into one CONOPS of forty-four enumerated elements: 4 PCE, 4 SEG, 15 SVC, 21 AST. Today the SOC steps forward with the threat picture, and you enumerate threats across kinetic, cyber, electronic-warfare, and naturally occurring domains against those same elements.

IN THE CLASSROOM

This module runs as slides, then a classroom of three parts. Have this workbook, the question set and the day's meeting minutes open. The classroom does this, in this order:

THE WORKBOOK	Chapter 3, The Gate, then Chapter 4, The Catalogue. Take two of the excluded reports in Chapter 4 and write the exclusion reason as the catalogue records it. Then take one threat that entered and name the one element it is anchored to.
THE QUESTION SET	Questions 1, 4 and 7 are worked and marked in the classroom, each answer given with its reason. The rest are yours.
LESSONS LEARNED FROM THE MINUTES	Open the Day 2 minutes at the decision register and read Item 01, How the day was going to be run, agreed before any evidence was opened; Item 02, What counts as a threat we can act on; Item 04, Threat 01, compromise of command authority. The rule, the objection and the decision are taken from the record, not from the slide.

What arrives. A credible peer-reviewed research report on command-and-control threats, prepared against the platform you defend, lands on your desk. It walks the three operational enclaves (Space, Link, Ground), names the potential attacks each one currently faces, and lists the platform subsystems each attack would touch. More sources follow: the four external types from the morning briefing and the SOC's own hunting.

What you build. Four `AN-THR` elements: 2 Ground, 1 Link, 1 Space, every one with a TOE naming specific `SVC` and `AST` elements from yesterday's decomposition. No free-floating threats; nothing in the set that the platform is not actually exposed to. Chapter 4 prints the whole deliverable.

Why it matters to the mandates. The catalogue feeds the risk analysis the NIS2 Directive requires under Article 21(2)(a), across the command-and-control scope Executive Order 14144 protects, and because every element is written in the shared five-field form it is ready for the policy-gated contribution to the Space ISAC Exchange the organization's own mandate requires.

What tomorrow consumes. Day 3, Converged Detection Engineering, picks up each threat's TOE and walks it into concrete attack paths and the data and signal sources that would reveal each

step. A clean TOE today is a traceable path tomorrow: that is why the targeting step is the strictest of the six.

CH 2 WHERE THREAT INTELLIGENCE COMES FROM

> "Threat intel comes from a range of sources."
>
> DEEP DIVE DECK, SLIDES 4 AND 5 · DAY-2 SOC BRIEFING: INTEL SOURCE TYPES · STARTING POINT

The sourcing discipline. The exact mix of sources depends on the organization; what matters is that every threat is sourced and cited to where it came from, which moves the organization off assumption-based defense and onto intelligence-driven, resilient operations. An uncited threat cannot be re-checked when the intelligence changes, cannot be weighed against a peer's report of the same activity, and cannot be defended in the risk analysis the mandates require. The source is what makes a threat citable and reproducible, which is why identifying it is step 02 of the method, before the threat even has an ordinal.

THE FIVE SOURCE TYPES

TYPE	ORIGIN	WHAT IT LOOKS LIKE	CITE AS
TI provider	External	A commercial threat-intel vendor feed, or the Space ISAC feed the organization's own mandate requires. That channel is live; consuming from it starts today.	vendor + advisory ID
Government	External	A federal advisory or attribution bulletin from a national cyber-defense agency.	agency + advisory reference
Peer-shared	External	An adversary profile or incident retrospective shared by another satellite operator.	peer org + retro ID
OSINT (open-source intelligence)	External	An open-source publication, peer-reviewed paper, or public threat database. Today's working source is a peer-reviewed study of the weaknesses that threaten satellite command and control.	publication + URL or DOI
Internal hunt	Internal	A threat your own Security Operations Center surfaced through in-house hunting on this platform.	internal hunt ticket + analyst

HOW TO READ AN INCOMING REPORT

Intel arrives. A credible research report on command-and-control threats lands on your desk: segment-by-segment narrative of potential attacks, with the subsystems each one would touch.

You assess it. Review each segment's narrative for the situational picture, then walk its table of potential attacks. Plausibility is judged with the departments: Satellite Operations weighs each threat against the platform as flown, Satellite Design & Engineering against the platform as built, the

same two views that produced the CONOPS. Match each plausible threat to the platform subsystems named in the report; set aside threats that have nothing to match.

> **Working rule.** The report proposes; the departments dispose. A finding enters the working list only after the two platform views agree it is plausible against this platform as flown and as built, and only if its targets resolve to enumerated CONOPS elements.

PART II

The method and the catalogue

The process for producing one entry correctly, then the complete catalogue this day produced, entry by entry.

CH 3 THE AN-THR METHOD

> "Six steps, fixed order: every threat's TOE naming the exact element it targets."
>
> DEEP DIVE DECK, SLIDE 8 · AN-THR ENUMERATION PROCESS

Inputs. Yesterday's decomposition: 44 ETENs of the telecommand path, the command-and-control scope Executive Order 14144 and the NIS2 Directive set. Threat intelligence cited to the five source types from the morning briefing; today it is the peer-reviewed research report in hand, read against the morning's composite OSINT assessment.

Output. One AN-THR ETEN per threat you add to the enumerated set, each one targeting one decomposed element via its TOE.

The Gate. "A threat targeting one enumerated platform element." THR admits an adversarial capability or intent. If it names an actor with no element it would act on, it is intelligence, not an enumerated threat: name the TOE or leave it out. A threat with no target stays out of the enumerated set until either the target is decomposed or the threat is scoped out.

THE SIX STEPS, FIXED ORDER

1. **Enumerate the LAYER.** LAYER = `AN`, fixed. These are not platform parts, they are the SOC's findings about the platform, so every element sits in the Analytic Layer. Unlike a PCE, SEG, SVC, or AST, an analytic element has no parent; it attaches to structure through its target field instead.
2. **Identify the source.** Threat-intelligence provider, government attribution statement, peer-shared adversary profile, OSINT, or the SOC's own hunting. Recorded on every entry so every downstream reader knows the intelligence basis.
3. **Set the TAG to THR.** The element is a Threat within the Analytic Layer (`AN-THR`): an adversarial capability the organization must enumerate, not an attack path or a detection.
4. **Assign the ORDINAL.** Two digits starting at `00`, one ordinal per distinct threat. Ordinals are issued by the Security Operations Center's running threat register, one sequence across every threat the organization tracks. This register opened this week, so today's four threats are the first four it has issued and they read 00, 01, 02, 03. They will not stay consecutive: as the register takes in threats outside this scope, and as entries retire, gaps appear, and the rule that makes those gaps safe is that an ordinal is issued once and never renumbered.
5. **Enumerate the TOE.** Target of Exploitation: **one** enumerated platform element the source states the threat is directed against. Concrete, not broad, at the lowest layer the source resolves. One element, not a list: the blast radius is real but it is description, not target. This element is the shared anchor of the chain, so the attack path on Day 3, the detection signature on Day 4 and the resilience measure on Day 5 all name it too.
6. **Write the DESCRIPTION.** Identify the threat actor or threat class, describe the assessed capability, and cite the source. This is the line a peer operator will read to act on the element.

Repeat for the next threat. When the enumerated set covers the three departments' priority threat list, the AN-THR enumeration is current. The TOE is the seam Functions 03 and 04 read next, so it must resolve to CONOPS elements that already exist.

CH 4 THE CATALOGUE, COMPLETE

> "Four threats enumerated: Ground 2, Link 1, Space 1, every TOE naming CONOPS elements."
>
> DEEP DIVE DECK, SLIDE 11 · DAY-2 ENUMERATED THREATS (ALL 4)

How to read each entry. The TOE holds **one** enumerated platform element, at the lowest layer the source resolves. That element is the seam Day 3 starts the attack path from, the element Day 4 watches, and the element Day 5 hardens, so one identifier carries the whole chain. Everything else the threat would touch on its way through is in the description, where it belongs: a target field with four elements in it cannot tell a reader which one the finding is actually about. The element's parents are not repeated here; they are read off the Day-1 CONOPS whenever they are needed. Source for every Day-2 entry: OSINT, the peer-reviewed research report read against the morning's composite assessment.

Watch the anchors accumulate. After each entry below you get the anchor figure: one row per threat, the threat on the left, its `TOE` field in the middle, and on the right the one platform element that field names, drawn in that element's own layer colour. By the end of the chapter the figure is the contextualized threat model: four findings, four anchors, four elements.

The analytic layer is drawn by a target edge, never by containment, because an AN element has no parent. Magenta is Layer 5; the box on the right keeps the colour of the structural layer it belongs to, and that colour change across the arrow is the anchor crossing layers. Every figure links to its vector source.

GROUND · 2 AN-THR

AN:THR:Threat:00 · SATCOM management-network intrusion

State-sponsored actor gains persistent access to the SATCOM management network and abuses the provider-to-customer trust relationship to push a malicious modem update (KA-SAT / AcidRain class).

FIELD	VALUE
ETEN	`AN:THR:Threat:00`
TOE	`SVC:CP:Control Plane:09`
What that element is	the ground command-acceptance service
Source	OSINT: composite assessment + research paper

The precedent, in full. KA-SAT / AcidRain is the anchor case for this class. In February 2022, an intrusion into a SATCOM provider's management network pushed a malicious modem update across the provider-to-customer trust relationship and disabled tens of thousands of terminals across Europe, including users far outside the intended target. The threat is not the modem, it is the management network's authority over everything downstream: which is why the TOE names `SVC:CP:Control Plane:09`, the ground command-acceptance service, and nothing else. The credential store and the patch pipeline are how the actor gets there and what the damage rides on, and both are in the description.

> **FIELD NOTE** The SOC drafted one ground-intrusion threat covering both the management network and the commanding consoles; Satellite Operations pushed back, because the two ride different Control Plane services and different credential stores. Ruling: two ordinals, 02 and 04, each with its own TOE. First win: Day 3 traces two distinct attack paths, each landing on the exact service and assets it abuses, instead of one blurred ground intrusion.

AN-THR ANCHORS · 1 OF 4 ENTRIES · AFTER AN:THR:THREAT:00

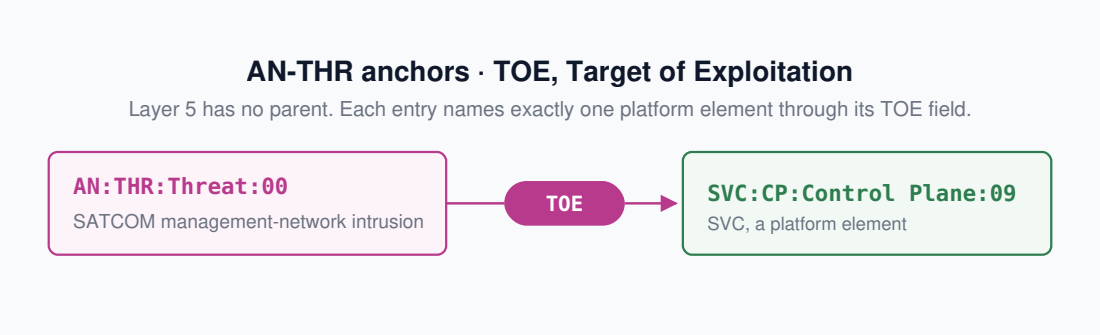

One finding, one anchor. The threat sits in Layer 5 on the left and names one element on the right, and the arrow through the `TOE` pill is the only link the two have. Nothing is nested, because an analytic element has no parent.

Why the box on the right is green. Green is the Service layer, the colour `SVC:CP:Control Plane:09` already wears in the Day-1 decomposition. Magenta is Layer 5 and nothing else. That colour change across the arrow is the anchor crossing layers, which is the one thing this figure exists to show.

AN:THR:Threat:01 · Command-authority compromise

Adversary with stolen or coerced command authority issues unauthorized telecommands to the on-orbit satellite.

FIELD	VALUE
ETEN	`AN:THR:Threat:01`
TOE	`AST:SW:Software:04`
What that element is	the operator console software
Source	OSINT: composite assessment + research paper

AN-THR ANCHORS · 2 OF 4 ENTRIES · AFTER AN:THR:THREAT:01

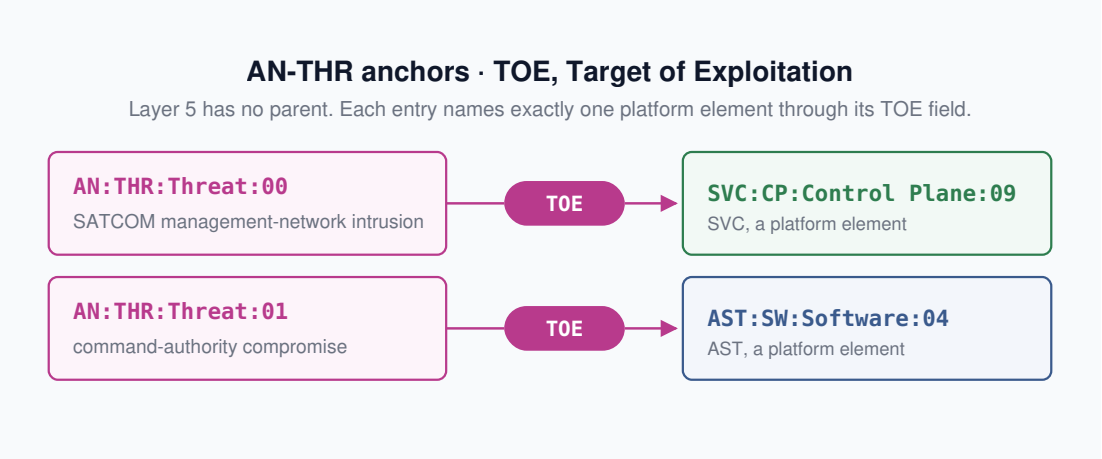

The second row anchors to an Asset, not a Service, so it draws in blue. Two findings, two different layers, one figure: the target field does not care which layer the element sits at, only that the source resolves to exactly one.

Why these two threats stayed apart. Both reach the ground commanding chain, and an earlier draft had them as one entry. They were split because they resolve to different elements, and an entry that resolves to two elements cannot be detected or hardened as one thing on Days 4 and 5.

LINK · 1 AN-THR

AN:THR:Threat:02 · Jamming campaign

RF actor sustains noise injection across uplink/downlink bands to deny communications.

FIELD	VALUE
ETEN	AN:THR:Threat:02
TOE	AST:SI:Signal:00
What that element is	the uplink and downlink waveforms
Source	OSINT: composite assessment + research paper

FIELD NOTE Satellite Design & Engineering questioned whether RF jamming belongs in a cyber threat set at all. The SOC applied the one test the framework gives, does the threat name one enumerated element, and it does: AST:SI:Signal:00, the uplink and downlink waveforms. Ruling: the TOE decides scope, not the attack medium, so jamming is enumerated. First win: the noise-floor telemetry Satellite Operations already watches is now pinned to AST:SI:Signal:00, ready to become a required signal source on Day 3 and the source a Day 4 signature fires on. A jamming threat is seen in the noise floor and never in a log, which is why Day 3 records data sources and signal sources as two separate fields.

AN-THR ANCHORS · 3 OF 4 ENTRIES · AFTER AN:THR:THREAT:02

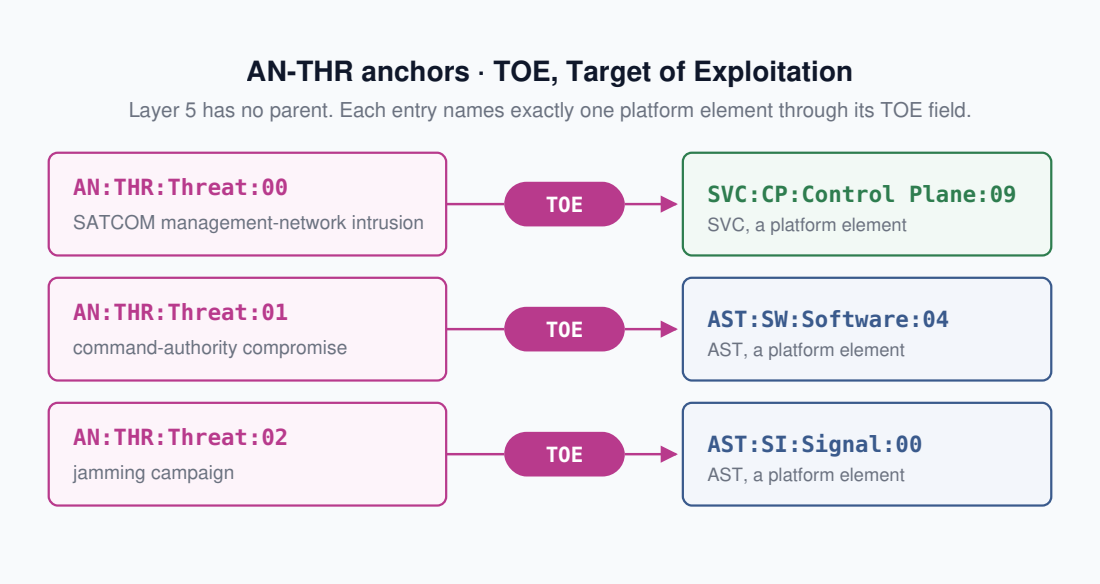

The link enters the catalogue. The waveform asset is the element the jammer acts on, so it is the TOE; the error-correction and tracking services that carry the waveform are reach, and they are in the description.

Why the services are not in the field. They would be four more identifiers with no way for a reader to tell which one the finding is about. The field answers one question, what is this threat directed at, and it answers it once.

SPACE · 1 AN-THR

AN:THR:Threat:03 · Malicious on-orbit firmware update

Adversary subverts the flight-software update supply chain and pushes a backdoored firmware update to the operational satellite over the command-and-update path.

FIELD	VALUE
ETEN	AN:THR:Threat:03
TOE	AST:FW:Firmware:01
What that element is	the on-board computer boot firmware
Source	OSINT: composite assessment + research paper

FIELD NOTE Satellite Design & Engineering read the firmware supply chain as a pre-launch, ground-side problem; the SOC's assessment showed the backdoored image executing in orbit, on the flight hardware. Ruling: a Space-segment threat, its TOE naming the on-board firmware assets and the update services that deliver them. First win: when Day 5 authors the measured-boot measure, its protected elements are already the right ones, the firmware in orbit, not only the pipeline on the ground.

AN-THR ANCHORS · 4 OF 4 ENTRIES · THE COMPLETE DAY-2 CONTEXTUALIZED THREAT MODEL

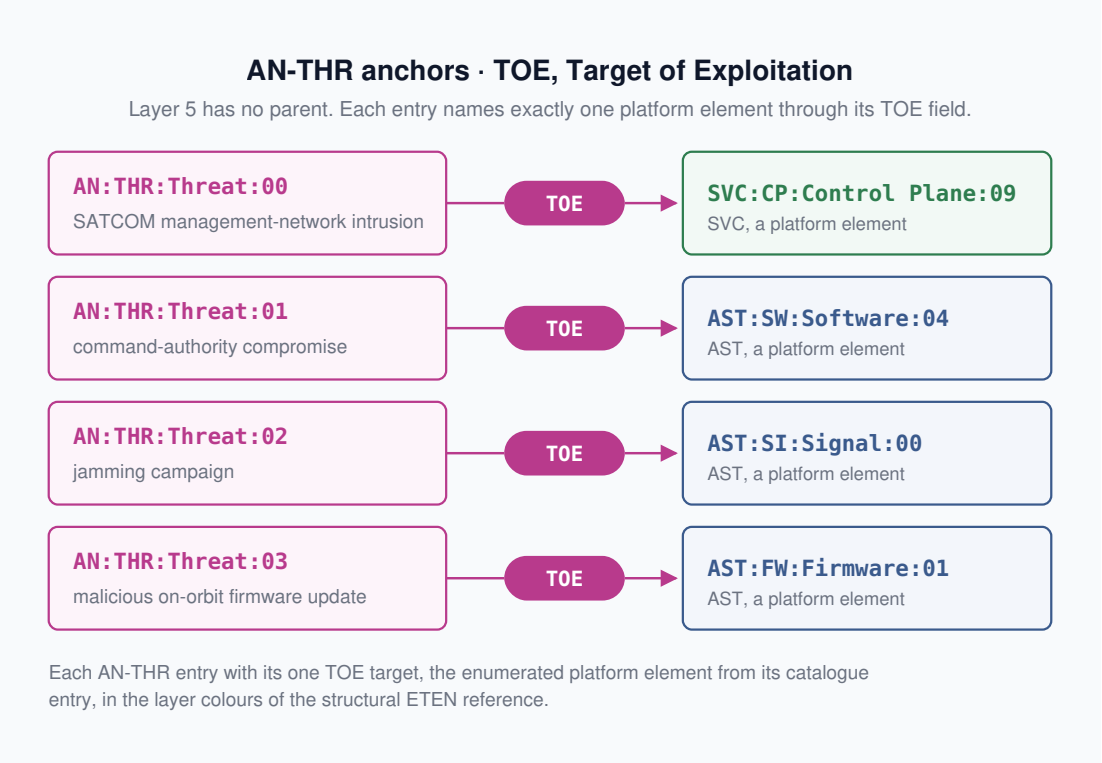

Four findings, four anchors, four elements: one service and three assets, reaching the ground, the link and the space segments. This figure is the deliverable Day 3 consumes, and every attack path it draws starts at one of these four boxes.

What four anchors do and do not claim. They do not claim the rest of the platform is safe. They claim that four threats, enumerated in one day from the evidence available, resolve to these four elements. The boundary is stated so a reader can hold it, and the register is built so the fifth entry costs nothing to add.

> **Scope note.** The User segment (`SEG-US-User`) is a segment type in the METEORSTORM data model, but it is out of scope on this platform's Day-1 CONOPS, so it is not enumerated and carries no AN-THR threats here.

PART III

Putting it to work

Who reads what you wrote, and how it leaves the room as machine-readable intelligence rather than a document a partner has to re-key.

CH 5 WHO READS THE CATALOGUE

> "One shared threat picture: three departments acting on the same enumerated set."
>
> DEEP DIVE DECK, SLIDES 11, 13 AND 17 · DAY-2 ENUMERATED THREATS (ALL 4) · PRESENT THE THREAT CATALOGUE · DAY 2 COMPLETE

The three readings. All three departments pick this set up in the same form. The Security Operations Center turns each element into a hunt priority against its TOE elements. Satellite Operations reads TOE to know which active subsystems carry which threats during operations. Satellite Design & Engineering reads TOE to plan hardening at the exact elements being targeted.

Three translations become one form. You can now stand in front of Security Operations, Satellite Operations, and Satellite Design & Engineering with one shared threat picture. You do not have to be the threat expert for every domain; you convene them and give them one contextualized picture to act on. Before this engagement that reading took three translations; now it takes one shared form.

The sharing seam. Because every element is written in the shared five-field form, the set is ready for the policy-gated contribution to Space ISAC that the organization's own mandate requires: comparable, citable, and machine-readable. Chapter 6 shows the tag that carries it.

PRESENTING IT BACK

The catalogue is enumerated; now you present it. You walk the three departments through the threat picture the same way you walked them through the CONOPS: the whole platform first, then each segment exploded with its threats targeting its own services and assets, then the full picture. The work is contiguous: CONOPS grounds threats, and what the departments confirm in this presentation is what tomorrow's work builds on.

CH 6 TAG AND SHARE

> “Your work ships as machine tags the whole community can read.”
>
> DEEP DIVE DECK, SLIDE 15 · TAG THE THREATS

The METEORSTORM taxonomy the Security Operations Center deployed on Day 1 is published openly as a MISP machinetag namespace. Tagging the four threats with the Analytic-layer threat entry turns today’s intel into machine-readable threat intelligence instead of a document a partner must re-key:

```
TAG meteorstorm:AN="AN-THR"
```

Each tag rides on the structural element the threat targets, so a peer operator reads your threat intel in the shared data model, with no translation. What a peer receives is exactly what chapter 4 prints: the five-field identifier, the description with its source citation, and the TOE naming enumerated elements, comparable against their own platform’s decomposition of the same element types.

Module 02 Contextualized Threat Modeling Workbook · Full Spectrum Space Cybersecurity Professional

4 AN-THR · Ground 2 + Link 1 + Space 1 · 17 of 44 elements targeted