FULL SPECTRUM SPACE CYBERSECURITY PROFESSIONAL

GUIDE · STUDENT EDITION

The Pentagon of Pain

Detect, disrupt, deter space threats. The mindset that judges security work by what it costs the adversary, its five mastery areas, and the exercises that make you apply it. Stands alone, inside or outside the course.

COPYRIGHT AND LEGAL NOTICES

PUBLISHER	D2 Team CORP, a 501(c)(3) nonprofit
PROGRAM	Full Spectrum Space Cybersecurity Professional
DOCUMENT	The Pentagon of Pain, guide, student edition
EDITION	Student edition · revised 2026-09-24
CONTACT	www.d2teamcorp.org

STUDENT EDITION

The teaching chapters and references. The exercises, with their answers, are the companion question set.

CONTENTS

CH 1 THE PENTAGON OF PAIN IN ONE PAGE

Space systems have become part of critical infrastructure, and so they have become inflection points for hybrid conflict. Defending them is no longer a matter of isolated expertise. It is a matter of shared architectural clarity, cross-domain insight and adversary awareness. Yet too often, space and cyber professionals operate in parallel rather than in sync, leaving exploitable seams between engineering, operations and security.

The question. Most security programs judge their work by what they have put in place. The Pentagon of Pain changes the question to what the adversary now has to pay. It is a way of judging, not a set of steps. It gives the engineers, the operators and the security analysts who defend one platform a single test for every hour and every dollar: does this raise the adversary's cost?

The five mastery areas. The Pentagon names five areas where mastery provably raises an adversary's cost, complexity and risk. Each one is a judgment, and each is stated as what the adversary suffers when the organization has mastered it.

#	MASTERY AREA	THE ADVERSARY SUFFERS WHEN
01	Master Decomposition	"The adversary suffers when you know your platform better than they ever can."
02	Master Contextualized Threat Modeling	"The adversary suffers when every strike they imagine is already prepared for."
03	Master Converged Detection Engineering	"The adversary suffers when they cannot hide, and every move is seen."
04	Master Exposure Management	"The adversary suffers when every path they take ends in a trap."
05	Master Adversary Management	"The adversary suffers when their plans are known, broken, and turned against them."

Mindset and method. The Pentagon of Pain is a mindset, not a framework. It does not tell an organization what to produce; it tells the organization how to judge what it produces. The method is the METEORSTORM cyber resilience framework, whose five functions, Concept of Operations, Contextualized Threat Modeling, Converged Detection Engineering, Incident Response Preparation and Adversary Management, produce the artifacts the judgment is applied to. Transform the mindset, then equip with the framework. A mindset with no framework is a poster. A framework with no mindset is compliance work with better filing. This guide carries the mindset. Where it names a function, it names it exactly as the framework defines it, so that the two never contradict each other whether you read this document inside the course or on its own.

Tradecraft, which is what this develops. A control is bought, a tool is installed, a policy is signed. None of those is what the Pentagon of Pain sets out to build. What the five mastery areas develop

is TRADECRAFT: the practiced methods and judgment that the engineers, operators and analysts defending one platform apply to it, learned by doing the work repeatedly rather than by reading about it once. Tradecraft is what a team can actually do when it sits down in front of the platform, and it is the first thing to decay when it is not practiced. The word is borrowed from the same discipline as the rest of this model, where tradecraft names the practiced craft of the analyst rather than the equipment they hold.

Why tradecraft is the thing to develop. An organization can hold every artifact the framework produces and still have no tradecraft, because artifacts record what was decided and tradecraft is the ability to decide it again, faster, against something the catalogue has not seen. Tools are replaced, vendors change, a platform is redesigned, and none of that removes tradecraft from the people who hold it. That is why each mastery chapter ends with something to do rather than something to file, and why the judgment is stated as five standing habits rather than five deliverables. The framework gives the tradecraft something to be practiced on; the mindset says whether the practice is working.

What this guide is. This is the Pentagon of Pain in full: the five mastery areas, the doctrine each is drawn from, the organizational argument, the worked case and the measure. It is a thought model, not an engineering specification, and it is the source document for the mindset rather than a summary of one held elsewhere. Where a mastery area could be mistaken for a function, this guide says which is which: a mastery area is a standing judgment, a function is work that produces an artifact, and the two carry different jobs and the same facts.

FIGURE 1 · THE PENTAGON OF PAIN

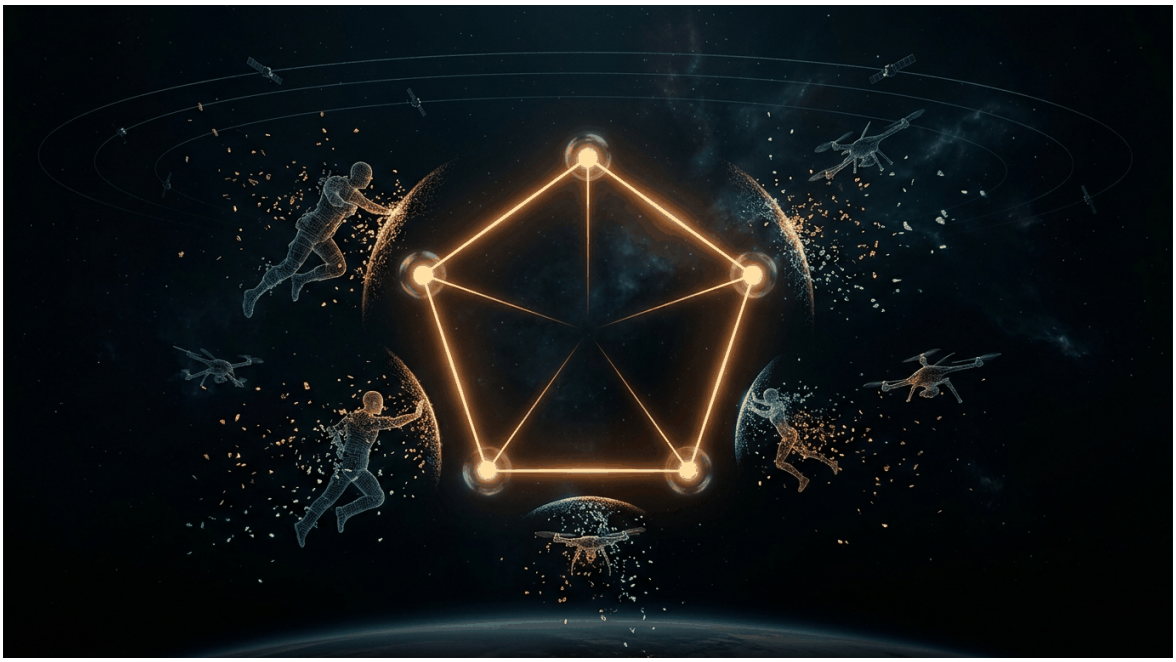

Read it as five points of one judgment, not five stages of a process. The numbering is the order in which the framework's functions run; the judgment applies to all five at once.

CH 2 DOCTRINAL FOUNDATIONS

The Pentagon of Pain is a defensive cyber operator's model, but its reasoning is borrowed from an established military discipline. This chapter states that borrowing plainly, so the reader can see where each mastery area comes from and can go to the source.

2.1 · TARGETING AS THE SOURCE DISCIPLINE

Air Force Doctrine Publication 3-60, *Targeting*, describes targeting as effects-based, interdisciplinary, anticipatory, systematic and integrated with other processes. It states that targeting is more than the selection of targets for destruction, that it draws on strategists, operators, intelligence personnel and legal advisers, and that it is inherently estimative: matching actions to targets requires anticipating outcomes under uncertainty. The joint targeting cycle in Joint Publication 3-60 runs through six phases: commander's objectives and intent; target development and prioritization; capabilities analysis; commander's decision and force assignment; mission planning and force execution; and combat assessment.

Two ideas in that doctrine matter most here. First, target development always approaches adversary capabilities from a systems perspective; a target's importance lies in its relationship to other targets within a system. Second, assessment measures whether desired effects were created and whether objectives were achieved, and planning for assessment begins before operations start. A defensive cyber operator who adopts both ideas stops asking whether a control exists and starts asking what its own platform looks like as a system, and whether the adversary actually paid a price.

Air Force Pamphlet 14-210, *USAF Intelligence Targeting Guide*, is the service's own how-to manual for that discipline, and it states the structure this guide borrows most directly. Its target system concept holds that a target is composed of components, and components are composed of elements; that each system is itself a component of a larger and more inclusive system; and that system components are interdependent, so a change in one component causes change in others. Read against a satellite platform rather than an enemy one, those three sentences are a parent chain and a blast radius. The same chapter directs the targeteer to focus not on the system or its components but on their activity, which is the distinction this framework draws between an asset and the service it carries.

That manual also lists environmental characteristics among the standing characteristics of a target, beside the functional, physical and mobility characteristics, and treats them as conditions within which the target exists that bear on how it is sensed and engaged. Naming the environment is therefore ordinary targeting practice of long standing. What is new is asking an operator to record it on their own platform, in the same place they record everything else.

Beneath the doctrine sits a standards layer that governs how the work is recorded. CJCSI 3370.01B, *Target Development Standards*, documents DoD policy and standards for target development and target intelligence. CJCSM 3314.01, *Intelligence Planning*, governs how

intelligence planning is conducted in support of a commander's plans. Both carry a release category of Limited, so this guide names them by designator and title and takes no text from them. They are cited because the reasoning here has a documented lineage, and a reader who holds the access can follow it down to the standards that govern the practice.

2.2 · READING DOCTRINE FROM THE DEFENSIVE CYBER OPERATOR'S SIDE

The Pentagon of Pain does not ask civilian operators to conduct targeting. It asks them to apply the targeteer's habits of mind to their own platform and to the adversary campaign directed at it. The table below is the spine of this guide. Each mastery area keeps its original name and intent; the doctrine column names the published process it is drawn from.

MASTERY AREA	DOCTRINAL ANCHOR	WHAT THE OPERATOR BORROWS
Master Decomposition	The target system concept in AFPAM 14-210, where a target is composed of components and components of elements, and each system is a component of a larger one; target system analysis and the five target characteristics (physical, functional, cognitive control and information, environmental, temporal) in AFDP 3-60 and JP 3-60; centre of gravity and critical vulnerability analysis in JP 5-0; target development standards in CJCSI 3370.01B.	Study your own platform as a target system: environments, segments, services, assets, and their control paths, and find the critical vulnerabilities before the adversary does.
Master Contextualized Threat Modeling	Joint intelligence preparation of the operational environment, especially evaluating the adversary and determining the adversary's most likely and most dangerous courses of action, in JP 2-01.3 and JP 2-0; intelligence planning in support of the commander's plans in CJCSM 3314.01; SPARTA and MITRE ATT&CK as the technique catalogues.	Anchor known adversary activity to the decomposed platform, under real mission and environmental conditions, rather than listing generic threats.
Master Converged Detection Engineering	The find, fix, track steps of the dynamic targeting kill chain (F2T2EA) in AFDP 3-60 and AFTTP 3-2.3; space domain awareness in JP 3-14; defensive cyberspace operations, internal defensive measures, in JP 3-12.	Trace the adversary's route through your own platform and name the data and the signals you would have to hold to see each step, treating space domain awareness and cyber detection as one sensing problem.
Master Exposure Management	Capabilities analysis and the temporal characteristic of targets in AFDP 3-60; positive control in Space Policy Directive-5; the resiliency goals of NIST SP 800-160 Volume 2.	Deny the adversary a stable pairing of capability to target: every path they take ends in a detection with an authorized response behind it.

Master Adversary Management	Combat assessment and reattack decisions in AFDP 3-60; intelligence-driven defense against adversary campaigns; defensive cyberspace operations, response actions, as the authorized boundary in JP 3-12; the instruments of national power (DIME).	Count what the adversary keeps reusing, harden it, assess whether the cost imposed changed their behavior, and route attribution to the authorities who can impose consequences.

2.3 · THE BOUNDARY THE MODEL RESPECTS

Doctrine draws a hard line between actions inside the defended network and actions outside it. JP 3-12 defines defensive cyberspace operations, internal defensive measures (DCO-IDM), as authorized defense actions within the defended portion of cyberspace, and defensive cyberspace operations, response actions (DCO-RA), as deliberate, authorized measures taken outside the defended network. Every action in this guide sits on the DCO-IDM side of that line. Where the guide speaks of consequences, escalation or the instruments of national power, it means that a civilian operator's contribution is attribution-quality evidence delivered through trusted channels, such as the Space Information Sharing and Analysis Center (Space ISAC), to the governments and coalitions that hold the authority to act.

2.4 · MEASURING PAIN THE WAY DOCTRINE MEASURES EFFECTS

Each mastery chapter ends with a Pain Index. The index borrows the distinction that AFDP 3-60 and JP 3-0 draw between measures of performance and measures of effectiveness. The left column is a measure of performance: the state of what the defending organization has produced. The Adversary Impact column is a measure of effectiveness: whether the desired effect on the adversary was created. The Pain score summarizes the second, not the first. An organization that has done a great deal of work but changed nothing about adversary behavior has performed without being effective, and the index is designed to say so.

> **The index is not a maturity model.** It is never self-scored. The left column names the state of a deliverable that can be inspected; the right column names what the adversary must now do that they did not have to do before. Cost must be concrete. "Raises adversary cost" is filler; name what the adversary must now do: hold access longer, burn a second capability, move where telemetry sees them, or beat three departments instead of one.

CH 3 ORGANIZATIONAL TRANSFORMATION

A framework adopted by one department is a filing improvement. A framework adopted by three departments that have never shared a model of the platform is a transformation. The Pentagon of Pain is aimed at the second outcome, and it is honest about the starting point: in most space organizations the people who know the platform, the people who know the adversary, and the people who built the platform sit in three different rooms and judge their work by three different standards.

3.1 · WHY DOCTRINE INSISTS ON THE TEAM

AFDP 3-60 states that targeting is interdisciplinary and that an effects-based approach is fundamentally a team effort: strategists and planners bring context, operators bring execution experience, intelligence personnel bring analysis of adversary strengths and vulnerabilities, and legal advisers bring the rules. No single specialty is allowed to run the cycle alone, because no single specialty holds enough of the picture. The defensive cyber operator's version of that team is built from three departments that already exist.

3.2 · THE THREE DEPARTMENTS AND THEIR SEAMS

Each department below is competent inside its own boundary. The gaps are not failures of skill; they are the seams between boundaries, and they are exactly where the Viasat attacker walked through.

DEPARTMENT	WHAT IT KNOWS	THE GAP	WHAT CHANGES
Satellite Operations (the satellite operations center)	Normal. Telemetry baselines, commanding procedure, manoeuvre windows, link conditions, contact schedules. The first to see an anomaly.	Reads every anomaly as an engineering fault or an environmental effect. Holds no adversary model, so an attack that looks like a fault is handled as a fault. Trusts the commanding path because it has always worked.	Becomes the owner of the temporal and environmental facts in the decomposition, the source of baselines for detection, and the authority on what a response action does to positive control.
Security Operations (the security operations center)	The adversary. Techniques, indicators, log analysis, incident handling, threat intelligence exchange. The first to know who is coming.	Sees enterprise logs, not spacecraft telemetry, RF behavior or orbital context. Cannot tell a command anomaly from a scheduled manoeuvre, so alerts lack mission context and are tuned down. Treats the ground segment as an IT estate rather than a control system for a vehicle.	Becomes the owner of threat modeling and of the adversary picture, and the co-author with Satellite Operations of every mission-aware detection and response.

Satellite Design and Engineering	The platform as built. Architecture, interfaces, flight software and its update path, bus and payload behavior, verification history. The only department that can change the design.	Security requirements arrive late and generic, so the architecture is verified as built, not as attacked. The update and management paths are designed for reliability, not for an adversary holding credentials. The decomposition exists as schematics, not as a threat-informed model others can use.	Becomes the owner of the platform description as a living, shared model, and the engineer of the resilience measures that survive rehearsal without endangering the vehicle.

3.3 · THE ARC OF THE CHANGE

Transformation is visible only if the starting point is stated. Organizations that adopt the Pentagon of Pain tend to move through four stages, and each department moves at its own pace.

1. **Compliance.** Work is judged by whether a control exists. The adversary is not mentioned. This is the baseline for all three departments.
2. **Friction.** A department's view is exposed as partial, usually by another department contradicting it. Operations discovers that its fault was an intrusion; Security discovers that its alert was a manoeuvre; Engineering discovers that its update path is a target system.
3. **Cost.** The first time a department asks what the adversary pays rather than what the auditor wants. This is the turn, and it is the point at which the Pain Index starts to mean something.
4. **Ownership.** The department applies adversary-cost reasoning to its own work without being asked. Engineering designs the seam out; Operations reports the anomaly to Security in the adversary's terms; Security writes the detection in the platform's terms.

Space Policy Directive-5 expects this kind of integration when it calls for cybersecurity-informed engineering across the lifecycle and for plans to retain or recover positive control; NIST IR 8401 expects it when it assigns Cybersecurity Framework outcomes to the ground segment's terminals and operations centers. Neither document can make three departments share a model. The mindset does that, and the chapters that follow name, for each mastery area, the seam it closes.

3.4 · THE TEAM THAT HOLDS THE SHARED MODEL

The team that holds the shared model is the SCOR team, for Space, Cybersecurity, Operations and Resilience. It is drawn from the three departments, with the mandate to hold the shared decomposition, to run threat modeling and detection engineering together, to approve response actions against positive control, and to maintain the adversary picture. Whether an organization stands it up as a working group or as a department of its own, its measure of success is the Pain Index: not the number of controls, but the observed change in adversary behavior.

> **The adversary pays when three departments hold one model of the platform.** Satellite Operations knows normal. Security Operations knows who is coming. Satellite Design and Engineering knows how the platform was built. The Pentagon of Pain is what happens when they judge their work by the same measure.

CH 4 MASTER DECOMPOSITION

> "The adversary suffers when you know your platform better than they ever can."
>
> MASTERY AREA 01 · JUDGED AGAINST THE FUNCTION CONCEPT OF OPERATIONS

EPIGRAPH "If you know the enemy and know yourself, you need not fear the result of a hundred battles." Sun Tzu, as popularly rendered.

4.1 · WHAT IS DECOMPOSITION?

Decomposition is the structured act of disassembling a complex system into granular, mission-relevant elements to establish a shared operational understanding across disciplines. In space cybersecurity, this goes well beyond system schematics. It becomes an alignment mechanism, enabling engineers, operators and cybersecurity teams to converge on a unified model of system behavior. Done properly, decomposition yields a map of the system's logic and dependencies that is not merely descriptive but anticipatory. It frames security within the actual behavior and mission function of the platform, so that defensive cyber operators perceive not isolated hardware but interacting flows of telemetry, command, timing and risk.

4.2 · DOCTRINAL GROUNDING

AFDP 3-60 states that target development always approaches adversary capabilities from a systems perspective, that a target system is a collection of assets directed to perform a function, and that a single target's importance lies in its relationship to other targets within a system. Decomposition is that same analysis performed by the owner on the owner's own platform. The targeteer's five target characteristics give the decomposition the questions to hold over every element:

- **Physical:** what the element is, as discernible by sensors and signatures.
- **Functional:** what the element does within the system, its activity level, and its importance to the mission.
- **Cognitive control and information:** how the element exercises control or processes information; doctrine notes that nearly every system has a central controlling function whose loss changes the whole system's behavior.
- **Environmental:** how the environment acts on the element and its surroundings.
- **Temporal:** when the element is vulnerable, and for how long.

JP 5-0 adds the centre of gravity method: identify critical capabilities, the critical requirements they depend on, and the critical vulnerabilities among those requirements. AFDP 3-60 notes that critical vulnerabilities are hard to tell apart from critical requirements and hard to translate into explicit

target sets, and that this translation is the foundation of a course of action. A decomposition that stops at a block diagram has not yet done this work.

4.3 · WHY IS IT NEEDED?

Without decomposition, security defaults to surface-level controls, obscuring latent risks at interface transitions, control handoffs and timing-sensitive subsystems. These blind spots, often overlooked in traditional documentation, create persistent pathways for adversaries to operate undetected. Falco et al. emphasize that the lack of technical standards for the secure design and operation of space systems has enabled scenarios in which permanent loss of satellite control occurs through under-mapped components and ambiguous system boundaries. NIST IR 8401 responds to the same gap for the ground segment by requiring organizations to classify the systems, processes and components of satellite command, control and payload operations before protecting them.

Effective decomposition closes these gaps by enabling teams to operate from a shared, mission-aligned architecture; anchoring security decisions in actual system behavior rather than intended function; constraining adversary movement by identifying and monitoring flow-dependent attack paths; supporting detection placement at high-risk interaction points; and increasing friction for AI-enhanced adversaries by surfacing edge conditions where stealth degrades.

4.4 · THE FUNCTION IT IS JUDGED AGAINST: CONCEPT OF OPERATIONS

The framework function that produces the decomposition is Concept of Operations. Its purpose is to decompose the platform tailored to your requirements. It works in four structural layers, top down, every child naming its parent, and the mastery area is judged against exactly that output:

- **Primary Capability Environment (PCE):** the operational zone in which a capability exists or is exercised. The root of the chain, and the layer that lets an operator deconflict a natural effect from an adversary using the environment as terrain.
- **Segment (SEG):** the enclaves of the system. Space, Link, Ground and User are among the published segment types.
- **Service (SVC):** the control and data jobs being done.
- **Asset (AST):** the concrete elements: hardware, firmware, software, data, signal. An asset names every service it carries, so the blast radius of losing it is honest.

Every element gets a name, a parent and a description in one shared nomenclature, and the decomposition is scoped to the mission objective rather than exhaustive by default: you state what is out and why. Existing artifacts, architecture diagrams, system design documents, mission CONOPS and asset inventories, are reused and labelled with the taxonomy rather than rebuilt from a blank page.

> **A distinction worth holding.** Decomposition is often taught as a set of groupings: space, ground, user, link, integration, with physical and logical assets catalogued beneath them. Judged against the function, the structure is the four layers above. Space, Ground, User and Link are segments within it, and integration is a property of the parent chain rather than a grouping of its own. The five targeting characteristics stay as the questions held over each element, not as the shape of the model.

4.5 · TAKE ACTION

Applied decomposition is not a postmortem task; it is a live, collaborative exercise conducted throughout the system lifecycle. It begins with mission objectives and proceeds through architectural elements to the interfaces where risk resides. AFDP 3-60 says the same of target system analysis: it should begin well in advance, draw on everything already known, and continue throughout. The cross-department team works not as compliance reviewers but as architecture stewards, tracing system flows, validating interface assumptions, and identifying where adversaries are most likely to hide or pivot. The decomposition becomes a living reference, the foundation for threat modeling, detection engineering, response preparation and resilience.

4.6 · CLOSING THE DEPARTMENTAL GAP

Satellite Design and Engineering already holds the decomposition, but as schematics written for verification. Satellite Operations holds the temporal and environmental facts, contact windows, manoeuvre phases and link conditions, but in procedures rather than in the model. Security Operations holds neither. The transformation is a single shared model that Engineering supplies, Operations annotates with time and environment, and Security reads to place detections. The adversary's cost rises the moment three departments can point at the same element and mean the same thing.

4.7 · PAIN INDEX

STATE OF THE DECOMPOSITION	ADVERSARY IMPACT	PAIN
Basic system diagrams exist but lack operational traceability or cross-functional ownership.	Attackers exploit undocumented interfaces and data paths, maintaining stealth and mobility.	1
The in-scope platform is enumerated across the four layers, every element with a parent and a description, agreed by the three departments.	Attackers lose stealth; a deviation from baseline on a named element raises a question all three departments can read.	2
Threats, attack paths, detections and resilience measures are anchored to the enumerated elements.	Lateral movement and command injection on the mapped platform are flagged and constrained against a named element.	3
The decomposition is re-run as a pass, extended where earlier exclusions said it should be, and the measures feed the next platform design.	Campaigns collapse mid-operation; attackers must retool and re-infiltrate under significantly greater risk each pass.	4

CH 5 MASTER CONTEXTUALIZED THREAT MODELING

> "The adversary suffers when every strike they imagine is already prepared for."
>
> MASTERY AREA 02 · JUDGED AGAINST THE FUNCTION CONTEXTUALIZED THREAT MODELING

EPIGRAPH "Victorious warriors win first and then go to war, while defeated warriors go to war first and then seek to win." Sun Tzu, as popularly rendered.

5.1 · WHAT IS CONTEXTUALIZED THREAT MODELING?

Threat modeling is a structured analytical discipline aimed at identifying how an adversary might compromise the mission of a space platform. Within the Pentagon of Pain it is not a checklist-driven or compliance-centric exercise. It is a strategic activity grounded in how your specific systems work and how real adversaries may behave. The process begins with targeted, scenario-aware questions. Which parts of the system would attract adversarial interest? Through which elements could they gain access or escalate control? How might they persist, move laterally or alter telemetry? What would be the operational consequences? Answering these demands more than threat enumeration. It calls for modeling how threats land on specific elements of the platform, which moves threat modeling into a predictive, adversary-facing discipline.

5.2 · DOCTRINAL GROUNDING

Joint intelligence preparation of the operational environment is the doctrinal home of this mastery area. JP 2-01.3 describes it as four steps: define the operational environment; describe its impact; evaluate the adversary; and determine the adversary's potential courses of action, in particular the most likely and the most dangerous. AFDP 3-60 states that this preparation should assist commanders in anticipating enemy intent and enable them to preempt enemy action, and that space, cyberspace and the electromagnetic spectrum should already be fully integrated into it.

Read from the defensive cyber operator's side: the decomposition is step one; the environments the platform operates in are step two; the adversary's known techniques, drawn from SPARTA for spacecraft and ground segments and from MITRE ATT&CK for enterprise and industrial control systems, are step three; and the adversary's route against this platform is step four. The framework records step three as the threat catalogue and step four as the attack path, in the next mastery area, so that each is anchored and each can be inspected.

5.3 · WHY IS IT NEEDED?

In contrast to static risk matrices, which often lack context specificity, threat modeling rooted in system behavior enables anticipatory defense. Falco et al. describe this strategic anticipation as the ability to create predictive friction: a condition where defensive cyber operators influence attacker decisions before those decisions produce mission impact. Doctrine calls targeting estimative and anticipatory for the same reason; it does not assume perfect knowledge, and it accepts that

uncertainty and friction still apply. Defensive cyber operators equipped with anchored threat models can redesign components into secure blocks that constrain adversarial manoeuvre, forcing attackers either to deviate from known pathways, exposing their presence, or to fail outright.

5.4 · THE FUNCTION IT IS JUDGED AGAINST: CONTEXTUALIZED THREAT MODELING

The framework function of the same name has the purpose of anchoring every threat to the platform you just decomposed. For every element, ask which threats actually apply to that element. The output is a threat catalogue where every entry names exactly one enumerated element as its target, at the lowest layer the source resolves, plus a record of what was excluded and why.

- **The gate.** A threat enters the catalogue only when it can name an enumerated target. When it cannot, it is excluded on the record with one of two reasons: the element is not enumerated yet, which is a reason for the next decomposition pass to run, or the platform has no such element. An exclusion is a deliverable. A catalogue with no exclusions has almost certainly forced something in.
- **Evidence.** Every entry is built on attributed adversary activity: threat-intelligence reporting, government attribution, or a peer-shared adversary profile. Confidence is stated with its reason, including where it is low. An entry with no source is not a threat; it is a worry.
- **The anchor starts the chain.** The element a threat names is what its attack path anchors to, what its detection watches and what its resilience measure hardens. Choosing it carelessly misdirects three later functions.

> **A distinction worth holding.** Unintentional disruptions such as radiation effects and atmospheric attenuation are often counted as threat vectors beside adversary action. Judged against the function, a threat is adversary activity; the environment enters the model one layer down, as the Primary Capability Environment, and that placement is what lets an operator ask whether a symptom was the environment or an adversary using it as terrain. The underlying point stands: an adversary who exploits beam divergence during attenuation is a threat, anchored to the link element it lands on, and the attenuation itself is the environment that element lives in.

5.5 · TAKE ACTION

To translate threat modeling into practice, embed it in design, detection and operational workflows. Initiate with decomposition outputs; apply structured technique catalogues, SPARTA, MITRE ATT&CK and ATT&CK for ICS, so that the catalogue carries campaign realism; involve the three departments so that threats are evaluated across technical and mission-execution contexts; and state, for each priority threat, its source, its confidence and the single element it targets, so that the attack path, the detection and the resilience measure have something concrete to serve. Adversary emulation is the defensive cyber operator's rehearsal, and AFDP 3-60 is explicit that well-organized, well-rehearsed procedures are what make time-compressed decisions possible.

5.6 · CLOSING THE DEPARTMENTAL GAP

Security Operations knows the adversary's techniques but models them against an enterprise network it understands, not against a spacecraft and ground segment it does not. Satellite Operations knows which windows matter and which faults recur, but has no vocabulary for an adversary who chooses those windows on purpose. Engineering knows which interfaces were never meant to be reachable. The transformation is a catalogue that Security leads, that Operations grounds in real mission conditions, and that Engineering answers with the element the threat actually lands on. A threat can only be anchored when all three are in the room.

5.7 · PAIN INDEX

STATE OF THE THREAT CATALOGUE	ADVERSARY IMPACT	PAIN
Threat modeling is ad hoc and focused on listing possible threats without alignment to the decomposed platform.	Attackers face minimal disruption; threat models are too generic to anticipate or constrain actions.	1
Known threats are anchored by METEORSTORM contextualized threat modeling, one enumerated element each, with source and confidence stated and exclusions recorded.	Attackers must vary tools or timing; reused techniques and AI-generated steps land on a named element the defensive cyber operator has already read.	2
Every anchored threat carries an attack path, a detection and a resilience measure on the same element.	Adversary campaign variants become predictable and less effective; the strike they imagined is the one the platform prepared for.	3
The catalogue is refreshed as reporting arrives and the exclusions drive the next decomposition pass.	Automated campaigns break under real-time counterplay; adversaries are forced into improvisation.	4

CH 6 MASTER CONVERGED DETECTION ENGINEERING

> "The adversary suffers when they cannot hide, and every move is seen."
>
> MASTERY AREA 03 · JUDGED AGAINST THE FUNCTION CONVERGED DETECTION ENGINEERING

EPIGRAPH "What is of supreme importance in war is to attack the enemy's strategy." Sun Tzu, as popularly rendered.

6.1 · WHAT IS CONVERGED DETECTION ENGINEERING?

Converged detection engineering is the integration of space domain awareness with cybersecurity detection engineering, a fusion that enables defensive cyber operators to detect adversary behavior not just in cyberspace but across the full range of orbital and terrestrial mission environments. This convergence marks a shift from siloed telemetry monitoring or isolated log analysis to a system-aware, threat-informed capability grounded in mission flow, domain context and adversary behavior. Detection is not confined to cybersecurity indicators (failed logins, port scans) or to space domain telemetry (orbital drift, anomalous signal strength). The two are synthesized so that anomalies in orbital behavior, link-layer performance, control system telemetry and cyber-access patterns are read as interdependent indicators.

6.2 · DOCTRINAL GROUNDING

Dynamic targeting doctrine describes six steps: find, fix, track, target, engage and assess (F2T2EA), and AFDP 3-60 states that the chain applies equally to non-lethal and non-kinetic effects, including space and cyberspace operations. The first three steps are a detection problem. The defensive cyber operator must find an anomaly, fix it to a specific element of the decomposed platform, and track it across segments and time. Doctrine also describes three levels of combat identification: whether an entity is friendly, hostile or neutral; what kind of thing it is; and what its intent is. A detection program that can only say something is anomalous has reached level one.

JP 3-14 names space domain awareness as a joint space mission area, and JP 3-12 places detection inside defensive cyberspace operations, internal defensive measures. Converged detection engineering refuses to let those two remain separate programs, because AFDP 3-60 already requires that space, cyberspace and the electromagnetic spectrum be analyzed together. The Lockheed Martin intrusion kill chain gives the cyber side of the same idea: detection placed at each phase of an adversary's campaign, so that one missed indicator does not mean a missed intrusion.

6.3 · WHY IS IT NEEDED?

Traditional detection architectures treat space system components as discrete domains: mission operations in one toolset, cybersecurity in another, space domain awareness in a third. Modern threats do not respect these boundaries. Adversaries increasingly use cross-domain techniques,

injecting malicious commands under the cover of orbital manoeuvres, spoofing telemetry during weather-related disruptions, or synchronizing cyber intrusions with physical-domain stress. Falco et al. highlight that space systems remain vulnerable to permanent loss of control where detection is not tied to system behavior and mission context. Benchoubane et al. show that link-layer degradation induced by atmospheric or adversarial interference can provide attackers with stealthy entry points unless detection accounts for both environmental and threat-induced anomalies.

6.4 · THE FUNCTION IT IS JUDGED AGAINST: CONVERGED DETECTION ENGINEERING

The framework function of the same name has the purpose of enumerating attack paths and the data and signals needed to detect each step. It has two deliverables, and the second is the one teams drop.

1. **The path catalogue.** For each threat, trace how an adversary could move through the platform to make it real, as an attack path anchored to the same element the threat names. The route, entry to objective, is recorded on the path; the objective need not be the anchor.
2. **The required-source inventory, carried on the path.** For each observable step, name the source a defensive cyber operator would have to be holding to see the adversary act, in two fields. A **required data source** is something received as a record: a log, an audit stream, an inventory, a telemetry feed. A **required signal source** is something measured in the physical layer: a noise floor, a carrier-to-noise ratio, a lock status, a received power level. Every record names the element it is observed on and a state: Available, Partial, Gap or unassessed. A missing source is a recorded finding, never a skipped step.

That second deliverable is what convergence means when it is written down. Two fields and not one, because they are two kinds of observable, the same line the data model draws between a data asset and a signal asset. On a space platform that line decides whether a path is detectable at all: jamming is seen in the noise floor and never in a log, so a jamming path with only a data field is a path nobody will ever see coming. The examples this guide uses, command execution out of expected orbital phase, telemetry replay aligned with manoeuvre windows, signal spoofing when strength or directionality deviates under environmental duress, are exactly the steps this inventory names a source for.

> **A distinction worth holding.** This mastery area is often read as writing custom detection logic and translating techniques into indicators of attack. Judged against the function, that is the next function's work: Incident Response Preparation writes the detection signature on the source this function named. Converged Detection Engineering decides what must be collected, and records honestly where nothing is being collected. Writing a rule before the path and the source inventory exist is the failure this function exists to prevent.

6.5 · TAKE ACTION

The cross-department team works alongside flight controllers, satellite engineers and defensive cyber operators to trace mission flows as attack paths and to instrument the mission-critical

pathways: command uplink, telemetry downlink, control timing and orbital dynamics interfaces. Behavioral baselining defines normal with temporal, spatial and cyber-physical fidelity, so that a required source can be judged Available rather than merely present. Scenario-driven rehearsal exposes blind spots, such as command spoofing during a low-Earth-orbit handover or anomaly masking during a space weather event, and each blind spot becomes a recorded Gap with an owner rather than a hope.

6.6 · CLOSING THE DEPARTMENTAL GAP

This is the mastery area where the seam is widest. Satellite Operations sees the telemetry and treats the anomaly as a fault. Security Operations sees the logs and cannot tell a command anomaly from a scheduled manoeuvre, so the alert is tuned down until it is silent. Neither sees what the other sees, and the adversary lives in the space between. The transformation is a path written jointly: Operations supplies what normal looks like, in time and in orbit; Security supplies what the adversary looks like, in technique; Engineering supplies where in the architecture the two can be observed together, and which interfaces the hardware actually permits. Converged detection engineering is not a tool purchase. It is three departments agreeing on one definition of suspicious and one list of what they would need to see it.

6.7 · PAIN INDEX

STATE OF THE PATH CATALOGUE AND SOURCE INVENTORY	ADVERSARY IMPACT	PAIN
Detection relies on default signatures or generic anomaly rules written before any path exists; no source inventory; space and cyber observables live in separate programs.	Attackers bypass detection by blending into noise or mimicking routine activity.	1
Each threat has a path anchored to its element, with the required data and signal sources named and their states recorded, gaps included.	Known techniques cross at least one step where a source is Available; simple replay or automation attempts begin to break down.	2
The gaps are being closed by the owning department, and the inventory is what the detection signatures are written against.	Tradecraft is exposed quickly; the attacker must sacrifice stealth, use untested techniques, or withdraw.	3
The inventory is re-assessed with every pass and the unassessed count is driven to zero.	Dwell time is crushed; attackers face immediate loss of cover and rapid exposure.	4

CH 7 MASTER EXPOSURE MANAGEMENT

> "The adversary suffers when every path they take ends in a trap."
>
> MASTERY AREA 04 · JUDGED AGAINST THE FUNCTION INCIDENT RESPONSE PREPARATION

EPIGRAPH "He who knows when he can fight and when he cannot will be victorious." Sun Tzu, as popularly rendered.

7.1 · WHAT IS EXPOSURE MANAGEMENT?

Exposure management is the continuous practice of identifying, minimizing and reconfiguring system interfaces, services and communication pathways to reduce the opportunities available to adversaries. Unlike traditional hardening, which fixes a static configuration at deployment, exposure management on a space platform demands adaptive, intelligence-informed control of the attack surface. This is not merely closing ports or disabling unused services. It means understanding which subsystems are discoverable, reachable or inferable from outside; mapping how those exposures vary with mission phase, orbital position and signal conditions; and answering adversary reconnaissance with a prepared response rather than an improvised one.

7.2 · DOCTRINAL GROUNDING

In the joint targeting cycle, capabilities analysis evaluates available capabilities against desired effects to find the pairing of capability to target that will work. The adversary performs that analysis against your platform. Exposure management is the defensive cyber operator's effort to deny a stable answer. AFDP 3-60's temporal characteristic makes the lever explicit: a target's vulnerability is described in terms of the time available, and the shorter and less predictable that window, the harder the pairing. Space Policy Directive-5 sets the boundary condition: whatever is changed, the operator must retain or be able to recover positive control of the space vehicle.

7.3 · WHY IS IT NEEDED?

In space operations, persistent access is often more valuable to adversaries than immediate exploitation. Attackers do not always act on first entry; they observe, wait, and act at moments of reduced vigilance such as manoeuvre phases, telemetry transitions or degraded link conditions. The Viasat attacker held access to the trusted management segment and then executed legitimate management commands at the hour of greatest effect. Without a prepared, authorized response, such footholds persist undetected, and a defense configured at commissioning is insufficient where adversaries evolve at machine speed.

7.4 · THE FUNCTION IT IS JUDGED AGAINST: INCIDENT RESPONSE PREPARATION

This is the one mastery area whose name differs from the function it is judged against, and the function is not renamed to match. The framework's fourth function is Incident Response

Preparation. Its purpose is to write the detection signatures and the response playbooks. It consumes the paths and the source inventory from the previous function and produces, for each path, one signature and one playbook.

- **The signature.** A detection that fires on the source the path named, anchored to the same element as the path and the threat. Every signature is written in an open, vendor-neutral format that carries the platform element and the playbook it triggers, so that a signature written by one operator deploys in another operator's stack without translation. A bare vendor query is not a signature.
- **The playbook.** The approved actions the organization is already authorized to take when the signature fires, across assess, contain, recover and report, each naming who decides and who executes and at what level of authority. A signature ships with its authorized response, or it does not ship.

Read the mastery area against that output. Exposure is the path the adversary counts on; the trap at the end of it is a signature on a named element with an authorized response behind it. One observation makes the bridge: a declared blind spot is itself an exposure, and the honest gap that all three departments have agreed to close is worth more than an undocumented control that one of them applied alone. The gaps the previous function recorded are the exposures this function closes, either with a signature on a newly collected source or with a compensating action in the playbook until the source exists. Every contain and recover action is bounded by positive control, which is why the actions that touch the vehicle are executed by Satellite Operations and why the person who decides never executes.

> **A distinction worth holding.** This mastery area is often described through the resiliency techniques of NIST SP 800-160 Volume 2, adaptive response, dynamic positioning, deception and unpredictability, and through reconfiguring exposed components mid-campaign. Those techniques are real. Judged against the functions, a change to the platform that outlives the incident is a resilience measure, which is the next function's deliverable and serves one of that publication's four goals. What this function delivers is the prepared response: the signature that sees the path being taken and the authorized actions that end it.

7.5 · TAKE ACTION

To disrupt modern campaigns, exposure must be actively closed, not passively monitored. The cross-department team writes, for every path, the signature on its required source and the playbook that runs when it fires; validates both through operational rehearsal to ensure continuity and resilience; and treats every recorded gap in the source inventory as an exposure with a named owner and a compensating action until it is closed. The objective is not to eliminate all exposure; that is neither feasible nor mission aligned. It is to make sure that every exposure the adversary can find is one the defensive cyber operator has already written the response to.

7.6 · CLOSING THE DEPARTMENTAL GAP

Exposure management fails in the seam between Security, which wants a reachable interface closed now, and Satellite Operations, which knows that closing it during a contact could cost positive control of the vehicle. Engineering, which designed the interface for reliability, is often not consulted until the argument is over. The playbook is that argument settled in advance, with all three at the table: Security names the exposure and the adversary who would use it, Operations sets the window in which the action is safe and executes the actions that touch the vehicle, and Engineering builds the control so that it survives rehearsal. Who decides and who executes is written down before the signature fires, so that nobody has to invent a response, or go looking for someone with authority, at the moment of greatest effect.

7.7 · PAIN INDEX

STATE OF THE SIGNATURES AND PLAYBOOKS	ADVERSARY IMPACT	PAIN
Exposure is managed via static policy and periodic review; signatures are vendor-locked with no paired response; responses are invented on the day.	Attackers exploit long-lived weaknesses and a predictable architecture without time pressure.	1
Each path has a signature on its required source and a playbook naming who decides and who executes, but some sources are still Partial.	Automated tools and known exploits begin failing on the instrumented routes; attackers must adapt tools or pivot to a route not yet instrumented.	2
Every recorded gap has a compensating action in a playbook and an owner closing it.	Intrusions lose persistence quickly; the blind spot they counted on is now a hold, a review or a rate limit with a named executor.	3
Signatures and playbooks are rehearsed and refreshed as the source inventory changes.	Attackers are expelled mid-operation and forced to rebuild access repeatedly at escalating cost, inside positive control.	4

CH 8 MASTER ADVERSARY MANAGEMENT

> "The adversary suffers when their plans are known, broken, and turned against them."
>
> MASTERY AREA 05 · JUDGED AGAINST THE FUNCTION ADVERSARY MANAGEMENT

EPIGRAPH "Know the enemy's plans, and you will be able to determine which strategy will succeed and which will fail." Sun Tzu, as popularly rendered.

8.1 · WHAT IS ADVERSARY MANAGEMENT?

Adversary management is the deliberate, intelligence-driven practice of tracking, understanding and countering the specific actors targeting your system. It extends beyond tools, exploits or indicator signatures to examine the strategic intent, campaign infrastructure and behavioral fingerprints of real threat actors. Rather than asking what vulnerability was exploited, adversary management asks what campaign is being executed, what it keeps reusing, and how to take that away. This requires a mindset realignment from incident response to campaign interdiction, and it positions attribution not as a compliance task but as an input to deterrence.

8.2 · DOCTRINAL GROUNDING

Two pieces of doctrine define this mastery area. The first is assessment. AFDP 3-60 states that assessment measures whether desired effects are created, whether objectives are achieved, and what next steps are required; that it is more than battle damage and more than an after-action intelligence function; and that planning for it begins before operations start. If the desired effects are not achieved, targets may be re-planned or different targets selected. Adversary management is that loop run against the adversary campaign: did the friction imposed by the other four mastery areas change what the adversary did, and if not, what changes next?

The second is the boundary. JP 3-12 reserves response actions outside the defended network to authorized forces. A civilian operator's adversary management therefore ends at attribution-quality evidence and its delivery. AFDP 3-60 describes the diplomatic, informational, military and economic instruments of national power as the means of creating effects beyond destruction; those instruments belong to governments, and the Viasat case shows the route: operator forensics, national agency analysis, coordinated public attribution, and sanctions. The intelligence-driven defense literature supplies the working methods: kill chain analysis of campaigns and the Diamond Model's linking of adversary, capability, infrastructure and victim.

8.3 · WHY IS IT NEEDED?

Most organizations halt at containment. Once a breach is addressed and services restored, the operation is considered complete. This passivity cedes strategic ground. Adversaries evolve, using each intrusion to refine techniques, harden infrastructure and expand reach. Without hardening what they reuse and without attribution, tactics improve unchecked, infrastructure remains viable

and unburned, and campaigns scale with impunity. Falco et al. emphasize that modern defense must move beyond patch-and-repair cycles and interdict campaigns by correlating technical artifacts with behavioral signatures and known actor profiles.

8.4 · THE FUNCTION IT IS JUDGED AGAINST: ADVERSARY MANAGEMENT

The framework function of the same name has the purpose of shrinking the attack surface the threats keep using. It consumes the threats, the paths and the signatures, and produces a portfolio of resilience measures.

1. **Count the recurrence.** Find the elements that appear across many threats and many paths. Count it from the catalogues, never from impression. A measure on a recurring element closes several routes at once, and that ratio is the argument you take to the department that has to build it.
2. **Build the measure.** Each measure protects one element, the same one the threat it counters names; serves one of the four goals of NIST SP 800-160 Volume 2, Anticipate, Withstand, Recover or Adapt; and has one owner and a cadence the technology actually supports.
3. **Walk the coverage.** Measures are enumerated before they are exercised, so a reviewer can walk every attack path and check that each is either covered or carries a consciously accepted residual risk with an owner and a date. For a vehicle already on orbit many measures cannot be applied at all, and the honest output is that the entry becomes input to the next generation of platform design. That is a real deliverable, not a failure.

This is how the mastery area's question, did the cost imposed change what the adversary does next, becomes a question with evidence. The recurrence count says what the campaign keeps reusing. The coverage walk says whether each route is now closed, or carries an owned residual risk, or neither, and neither is the finding.

> **A distinction worth holding.** This mastery area is often written as actor profiling, attribution and campaign interdiction, with trusted intelligence exchange through Space ISAC among its actions. All of that is kept. Judged against the functions, the attribution evidence and the sharing go out through the report stage of the playbook, which is the previous function's deliverable, and the function of Adversary Management itself is the resilience portfolio and the coverage walk. The two together are the whole picture: the campaign is tracked, what it reuses is taken away, and the evidence reaches the parties able to impose consequences.

8.5 · TAKE ACTION

The cross-department team maintains living adversary profiles enriched with the elements each actor has been observed to land on; counts recurrence from the catalogues after every pass; builds one measure per recurring element with a goal, an owner and a cadence; records every accepted residual risk with an owner and a date; shares attribution-quality evidence through trusted communities such as Space ISAC and allied government channels, through the report stage of the playbook; and feeds the portfolio into the next platform design. At its highest maturity, adversary

management shapes attacker decision-making. It introduces risk, imposes cost, and forces recalibration not just at the keyboard but across adversary planning tables.

8.6 · CLOSING THE DEPARTMENTAL GAP

Adversary management is the mastery area that shows whether the transformation took. If Satellite Operations still files the anomaly as a fault, there is no recurrence to count. If Engineering never learns which of its elements the adversary keeps landing on, the next design carries the same seam. If Security tracks the actor but cannot describe the effect in the platform's terms, the evidence it hands to Space ISAC or to a national agency is thin. The transformation is one portfolio: Engineering owns the rebuilds, Operations owns the procedures that carry a measure until the rebuild flies, and Security owns the evidence that the measure changed what the adversary did. A department that asks that question of its own work has reached the ownership stage.

8.7 · PAIN INDEX

STATE OF THE RESILIENCE PORTFOLIO	ADVERSARY IMPACT	PAIN
Response ends at containment; the same flaw stays exposed; no recurrence is counted; attribution is ad hoc and sharing happens after the fact if at all.	Attackers retain plausible deniability and can reuse infrastructure or tactics without consequence.	1
Recurrence is counted from the catalogues and each measure protects one element for one goal with an owner and a cadence; some sharing occurs after containment.	Attackers begin to lose the element they reused most, and suffer exposure across intelligence-sharing communities.	2
Every path is either covered or carries an owned, dated residual risk, and attribution-quality evidence goes out on a confirmed attack.	Campaigns become traceable; attacker playbooks degrade quickly, and the residual path has an expiry date.	3
The residual risks close in the next platform design and the portfolio and the evidence are portable across operators.	Adversaries are identified and tracked in near real time; reuse becomes a liability and deniability collapses.	4

CH 9 DETECT, DISRUPT, DETER: THE VIASAT KA-SAT ATTACK

9.1 · DESIGN TO DETECT, DETECT TO DISRUPT, DISRUPT TO DETER

Design to detect. Strengthening systems during the design phase must move beyond superficial hardening. Ben Yahia et al. show that securing components from the outset of the design process addresses vulnerabilities before deployment. Building the required sources into the architecture, tied to mission behavior and telemetry flows, reduces blind spots and establishes early friction against AI-enhanced threats. This is the same principle SPD-5 states for the whole sector: cybersecurity-informed engineering across the lifecycle, not bolted on afterwards.

Detect to disrupt. Space platforms generate telemetry that can provide early warning long before exploitation completes. Benchoubane et al. warn that atmospheric attenuation and beam misalignment can notably degrade secrecy capacity, increasing adversary interception capability. By holding the signal sources that show such deviations, defensive cyber operators surface attacker actions and break the kill chain before lateral movement succeeds. Doctrine makes the same point from the other side: a target's temporal characteristic describes its vulnerability to detection in terms of the time available, and comparing that window to information latency tells the staff when and how to act.

Disrupt to deter. Campaign correlation is what turns detections into actionable intelligence. Falco et al. observe that permanent loss of satellite control arises from interconnected failures across system segments, requiring attribution that is strategic as well as technical. By linking detected behaviors to threat actors and mission impacts, defensive cyber operators enable consequences. The Viasat KA-SAT attack is the proof: months of forensic work by the operator and by national agencies produced a coordinated public attribution to Russia on 10 May 2022, and sanctions followed the same week. Organizations such as Space ISAC exist to move that evidence, and the sharing itself undermines anonymity and erodes attacker adaptability.

9.2 · WHAT HAPPENED

On 24 February 2022, at 05:02 local time in Ukraine and about an hour before the Russian invasion began, a cyberattack disrupted the KA-SAT satellite broadband network operated by Viasat. Viasat's own investigation found a ground-based intrusion in which the attacker exploited a misconfiguration in a VPN appliance to gain remote access to the trusted management segment of the KA-SAT network, moved laterally to the segment used to manage and operate the network, and then used that access to execute legitimate, targeted management commands on a large number of residential modems simultaneously. Those commands overwrote key data in the modems' flash memory, leaving them unable to access the network. Several thousand customers in Ukraine and tens of thousands of fixed broadband customers elsewhere in Europe were affected, and spillover left 5,800 Enercon wind turbines in Germany unable to communicate for remote monitoring or control. Researchers identified the wiper as AcidRain, now catalogued in MITRE ATT&CK as

software S1125. On 10 May 2022 the United States, the United Kingdom, the European Union and allied governments publicly attributed the attack to Russia, and sanctions followed the same week.

- **Initial access:** misconfigured VPN appliance exploited for remote access to the trusted management segment.
- **Propagation:** lateral movement within the management network, then legitimate management commands issued at scale.
- **Payload:** AcidRain wiper overwriting modem flash memory.
- **Impact:** cross-national loss of connectivity and loss of remote monitoring for industrial assets.
- **Attribution:** public, coordinated attribution to Russia on 10 May 2022, followed by sanctions.

9.3 · READING THE ATTACK THROUGH DOCTRINE

In targeting terms the adversary did everything the joint targeting cycle describes. They set an objective (deny Ukrainian communications at H-hour), developed the target as a system (the management segment that could reach every modem), analyzed capabilities (legitimate management commands plus a wiper), chose the timing (the temporal characteristic), executed, and assessed. The defensive cyber operator's Pentagon of Pain is the mirror image, and the table below shows where each mastery area, judged against its function, would have raised the adversary's cost.

ATTACK VECTOR	MASTERY AREA, FUNCTION	WHAT THE FUNCTION WOULD HAVE PRODUCED	ADVERSARY PAIN
Management segment reachable from a VPN appliance	Master Decomposition, Concept of Operations	The management and update paths enumerated as segments, services and assets with parents, the VPN appliance identified as a critical requirement and its misconfiguration as a critical vulnerability (JP 5-0 method).	Predictable lateral paths constrained; the management-to-modem command path understood and guarded by three departments who name it the same way.
VPN appliance compromise	Master Contextualized Threat Modeling, Contextualized Threat Modeling	A threat for third-party and remote-access abuse, anchored to the management service it lands on, sourced and with confidence stated; the most dangerous course, mass destructive commands at a moment of geopolitical stress, named in advance.	Access paths flagged before use; the strike is the one the platform prepared for.
Mass management commands to modems	Master Converged Detection Engineering, Converged Detection Engineering	An attack path from the appliance to the modem fleet, with the management-command record and the fleet-wide lock status named as required sources and their states recorded.	The push is observed on a named source; where nothing was collected, the gap is a finding with an owner, not silence.

Lateral movement within the trusted segment	Master Exposure Management, Incident Response Preparation	A signature on management command volume, timing and scope, paired with a playbook whose contain actions, isolating the management interface and rate-limiting management actions, preserve positive control and name who decides and who executes.	Malicious command behavior fires early with a rehearsed response behind it; propagation halted before scale.
Attribution and consequence	Master Adversary Management, Adversary Management	Segmentation of the management and update infrastructure as a resilience measure on the recurring element, to Anticipate, with an owner; operator forensics delivered through Space ISAC and national agencies via the playbook's report stage.	The high-impact path removed; public attribution and sanctions on 10 May 2022; strategic cost imposed through the instruments of national power.

9.4 · TAKEAWAYS

The Viasat incident shows how attacks on space-ground hybrid infrastructure exploit overlooked architectural seams, particularly at vendor interfaces and control-update boundaries. The absence of decomposition, threat modeling and instrumented detection left management and update services exposed. Each of the five mastery areas could have raised detection fidelity, invalidated trusted pivot paths and increased adversary friction, and the attribution that did occur shows the fifth working at national scale. Read as an organizational case, the attack ran through the seams this guide describes. The management segment was an engineering design decision, its reachability was a security exposure, and the mass of legitimate commands at 05:02 was an operational anomaly. Each department would have recognized its own piece. The adversary succeeded because no one held all three.

CH 10 SUMMARY

The Pentagon of Pain is the proactive ability to impose time, resource and uncertainty costs on adversaries through systemic design alignment. That ability is tradecraft, and tradecraft is what the five mastery areas develop: not a set of documents an organization holds, but what its people can do with the platform in front of them. The five mastery areas work together to create operational friction that forces attackers to slow down, rethink their approach, or abandon their objectives. Each has a home in published doctrine, policy or research, and each is judged against one function of the framework.

MASTERY AREA	DEFINITION	ADVERSARY PAIN	JUDGED AGAINST
Master Decomposition	Knowing the platform as an enumerated system: environments, segments, services and assets, each with a parent, so that attack surfaces and behavioral chokepoints are named rather than assumed.	Adversary must invest more time and resources to discover undocumented pathways, losing stealth and meeting hardened chokepoints.	Concept of Operations
Master Contextualized Threat Modeling	Anchoring attributed adversary activity to the specific elements it targets, with source and confidence stated and exclusions recorded, instead of listing generic threats.	Adversary must constantly adapt, as reused tactics land on a named element the defensive cyber operator has already read.	Contextualized Threat Modeling
Master Converged Detection Engineering	Tracing the adversary's route through the platform and naming the data and signal sources needed to see each step, with every gap recorded.	Adversary dwell time is cut sharply; attempts to blend into noise or exploit phase changes are rapidly exposed.	Converged Detection Engineering
Master Exposure Management	Ending every path the adversary can take in a detection signature with an authorized response behind it, bounded by positive control.	Adversary finds fewer viable attack paths, and persistent footholds are disrupted mid-operation.	Incident Response Preparation
Master Adversary Management	Counting what the adversary keeps reusing, hardening it as resilience measures with owners and cadences, and routing attribution-quality evidence to the authorities who can impose consequences.	Adversary loses anonymity and operational continuity, facing exposure, tool and infrastructure loss, and real-world consequences.	Adversary Management

> **Transform the mindset. Then equip the team.** The Pentagon of Pain changes the question an organization asks about its own resilience work, from what controls exist to what the adversary now has to pay. The framework's five functions give that question a method to act on. The organization that acts on it is not a new department; it is Satellite Operations, Security Operations, and Satellite Design and Engineering holding one model of the platform, one picture of the adversary, and one standard for judging their work.

DEF DEFINITIONS

TERM	DEFINITION
Adversary	Any entity, nation-state, criminal organization, insider or automated system, that intentionally seeks to compromise, disrupt or degrade the mission, security or operational integrity of a space platform or its supporting infrastructure.
Adversary impact	The measurable effect that the defensive cyber operator's work imposes on an adversary's campaign, including increased operational cost, reduced stealth, forced retooling, campaign disruption, or complete mission failure for the attacker. In doctrinal terms, a measure of effectiveness.
Attack path	The route an adversary could take through the decomposed platform to realize a threat, anchored to the element the threat names, and carrying the required data and signal sources needed to see each step.
Center of gravity (COG)	In JP 5-0, a source of power that provides moral or physical strength, freedom of action, or will to act. COG analysis identifies critical capabilities, critical requirements and critical vulnerabilities.
Combat identification (CID)	In AFDP 3-60, the characterization of a detected entity at three levels: friendly, hostile or neutral; type of platform; and intent.
Defensive cyberspace operations (DCO)	In JP 3-12, passive and active cyberspace operations intended to preserve the ability to use friendly cyberspace capabilities and protect data, networks and designated systems. DCO-IDM are authorized actions within the defended portion of cyberspace; DCO-RA are authorized actions taken outside the defended network.
Detection signature	A detection written on a required source, anchored to one platform element, in an open format that names the playbook it triggers.
Exposure	The degree to which system interfaces, services and communication pathways are discoverable, reachable or inferable by adversaries. Exposure is dynamic, fluctuating with mission phase, orbital position, environmental conditions and operational workload. A declared blind spot is itself an exposure.
F2T2EA	Find, fix, track, target, engage and assess: the six steps of dynamic targeting, which AFDP 3-60 states apply equally to non-kinetic effects including space and cyberspace operations.
Joint intelligence preparation of the operational environment (JIPOE)	The analytical process in JP 2-01.3 that defines the operational environment, describes its impact, evaluates the adversary, and determines the adversary's most likely and most dangerous courses of action.

Mastery area	One of the five points of the Pentagon of Pain: a judgment about what the adversary suffers when the organization has mastered it. A mastery area is judged against the framework function that produces the relevant artifact; it is not itself a function.
Mindset	The way an organization judges its own security work. The Pentagon of Pain is the mindset that judges work by what it costs the adversary. It is distinct from, and does not replace, the framework that produces the work.
Pain score	A number from 1 to 4 that summarizes the Adversary Impact column of a Pain Index, never the state column and never a self-assessment.
Positive control	In SPD-5, the operator's ability to retain, or if necessary recover, command of the space vehicle; the boundary condition on every response action.
Required data source, required signal source	The two kinds of observable a defensive cyber operator must hold to see a step of an attack path: a record received (a log, an audit stream, a telemetry feed) and a quantity measured in the physical layer (a noise floor, a carrier-to-noise ratio, a lock status). Each carries a state: Available, Partial, Gap or unassessed.
Resilience measure	A change that protects one platform element toward one of the four NIST SP 800-160 Volume 2 goals, Anticipate, Withstand, Recover or Adapt, with an owner and a cadence; where it cannot be applied to the flying vehicle it is input to the next platform design and the path carries an owned, dated residual risk.
Response playbook	The approved actions an organization is already authorized to take when a signature fires, across assess, contain, recover and report, each naming who decides and who executes.
Secure blocks	Architectural components redesigned through decomposition and threat modeling to incorporate security controls, clear control logic and traceability. Secure blocks constrain adversary movement by forcing attackers to deviate from known pathways, exposing their presence, or failing to achieve objectives.
Target system analysis (TSA)	In AFDP 3-60, the systematic examination of a target system to understand how its components interact and how the system functions; the analysis a defensive cyber operator performs on their own platform in decomposition.
Threat	Attributed adversary activity anchored to exactly one enumerated platform element, with its source and confidence stated. A candidate that cannot name an element is excluded on the record. Environmental effects are recorded as the environment the element lives in, not as threats.
Tradecraft	The practiced methods and judgment a defensive cyber operator applies to their own platform, developed by doing the work repeatedly rather than by reading it once. Distinct from the tools, controls and artifacts an organization holds: tradecraft is what its engineers, operators and analysts can actually do with them, and it is what the five mastery areas set out to build.

REF REFERENCES

1. G. Falco et al., “Minimum Requirements for Space System Cybersecurity: Ensuring Cyber Access to Space,” 2024 IEEE 10th International Conference on Space Mission Challenges for Information Technology (SMC-IT), Mountain View, CA, USA, July 2024, pp. 78 to 88, doi 10.1109/SMC-IT61443.2024.00016.
2. O. Ben Yahia et al., “Securing Satellite Link Segment: A Secure-by-Component Design,” 2024 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE), Daytona Beach, FL, USA, December 2024, pp. 177 to 182, doi 10.1109/WiSEE61249.2024.10850060.
3. N. Benchoubane, O. Ben Yahia, W. Ferguson, G. Gür, S. Chakravarty, G. Falco and G. K. Kurt, “Securing Heterogeneous Network (HetNet) Communications for Wildfire Management: Mitigating the Effects of Adversarial and Environmental Threats,” IEEE Journal of Radio Frequency Identification, vol. 9, pp. 713 to 726, 2025, doi 10.1109/JRFID.2025.3601843.
4. Curtis E. LeMay Center for Doctrine Development and Education, Air Force Doctrine Publication 3-60, Targeting, U.S. Air Force, 1 May 2026.
5. Targeting Division, Headquarters 497th Intelligence Group, Air Intelligence Agency, Air Force Pamphlet 14-210, Intelligence, USAF Intelligence Targeting Guide, U.S. Air Force, 1 February 1998. Supersedes AFP 200-17, 23 June 1989, and AFP 200-18 Volumes I and II, 1 October 1990. The target system concept, the target characteristics table and the six-phase targeting process are drawn from Chapter 1.
6. Joint Chiefs of Staff, Joint Publication 3-60, Joint Targeting, 20 September 2024.
7. Joint Chiefs of Staff, CJCSI 3370.01B, Target Development Standards, Joint Staff J-2, 6 May 2016 (distribution limited; named here by designator and title, with no text drawn from it). Supersedes CJCSI 3370.01A, 17 October 2014, and consolidates the cancelled CJCSM 3370.01, Target Graphics Standards, and CJCSM 3375.01, Target Intelligence Data Standards.
8. Joint Chiefs of Staff, CJCSM 3314.01, Intelligence Planning, Joint Staff J-2 (distribution limited; named here by designator and title, with no text drawn from it).
9. Joint Chiefs of Staff, Joint Publication 3-0, Joint Campaigns and Operations, 18 June 2022.
10. Joint Chiefs of Staff, Joint Publication 5-0, Joint Planning, 1 December 2020.
11. Joint Chiefs of Staff, Joint Publication 2-0, Joint Intelligence, 26 May 2022.
12. Joint Chiefs of Staff, Joint Publication 2-01.3, Joint Intelligence Preparation of the Operational Environment, 21 May 2014.
13. Joint Chiefs of Staff, Joint Publication 3-12, Joint Cyberspace Operations, 19 December 2022 (distribution limited; the 8 June 2018 edition, Cyberspace Operations, is publicly available and carries the same DCO-IDM and DCO-RA definitions).
14. Joint Chiefs of Staff, Joint Publication 3-14, Joint Space Operations, 23 August 2023.

15. Air Land Sea Space Application Center, AFTTP 3-2.3, Multi-Service Tactics, Techniques, and Procedures for Dynamic Targeting, as cited in AFDP 3-60 for the F2T2EA kill chain.
16. R. Ross, V. Pillitteri, R. Graubart, D. Bodeau and R. McQuaid, Developing Cyber-Resilient Systems: A Systems Security Engineering Approach, NIST Special Publication 800-160, Vol. 2, Rev. 1, December 2021, doi 10.6028/NIST.SP.800-160v2r1.
17. National Institute of Standards and Technology, Satellite Ground Segment: Applying the Cybersecurity Framework to Satellite Command and Control, NIST IR 8401, December 2022, doi 10.6028/NIST.IR.8401.
18. National Institute of Standards and Technology, Introduction to Cybersecurity for Commercial Satellite Operations, NIST IR 8270, July 2023, doi 10.6028/NIST.IR.8270.
19. The White House, Space Policy Directive-5, Cybersecurity Principles for Space Systems, 4 September 2020.
20. The Aerospace Corporation, Space Attack Research and Tactic Analysis (SPARTA), sparta.aerospace.org.
21. MITRE Corporation, MITRE ATT&CK, Enterprise and ICS matrices, attack.mitre.org.
22. Viasat, “KA-SAT Network cyber attack overview,” 30 March 2022.
23. U.S. Department of State, “Attribution of Russia’s Malicious Cyber Activity Against Ukraine,” press statement, 10 May 2022; with concurrent statements by the United Kingdom, the European Union and partner governments.
24. E. M. Hutchins, M. J. Cloppert and R. M. Amin, “Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains,” Lockheed Martin Corporation, 2011.
25. S. Caltagirone, A. Pendergast and C. Betz, “The Diamond Model of Intrusion Analysis,” Technical Report ADA586960, 2013.
26. Sun Tzu, The Art of War, L. Giles, trans., 1910 (public domain); quotations as popularly rendered.
27. MITRE Corporation, “AcidRain, Software S1125,” MITRE ATT&CK.
28. MISP meteorstorm taxonomy (machinetag.json, version 2): the published METEORSTORM namespace, the source of the layer and element definitions the framework functions use.

AI software was used in preparing this guide for editorial quality control and for locating and cross-checking the doctrinal sources cited. Every doctrinal statement here was checked against the published text of the publication it cites, and the two distribution-limited publications are named by designator and title with no text drawn from them.
