ACDP

Actor-Centric Defensive Prioritization

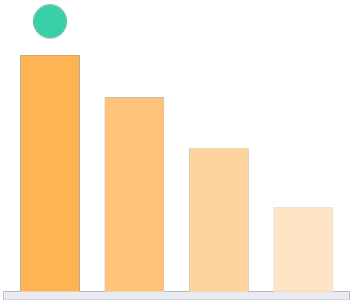

A Methodology for Aligning Defensive Investment with Adversary Intent

---

Technical Paper · **Malwarebox Research** · Version 1.0

Author: Robin Dost

Original blog post: blog.synapticsystems.de/actor-centric-defensive-prioritization-acdp

Published: April 2026

Abstract

Threat intelligence has matured rapidly in the last decade, producing rich descriptions of adversary capabilities, tooling and behavior. Defensive planning has not kept pace. Organizations frequently operate with sound technical controls and comprehensive best-practice catalogs, yet still fail to prevent strategic losses because their defensive investments are not aligned with the specific objectives of the adversaries they face.

This paper introduces Actor-Centric Defensive Prioritization (ACDP), a methodology that sits one level above existing frameworks such as MITRE ATT&CK and NIST. ACDP uses threat intelligence outputs as inputs and produces a reasoned ordering of defensive actions based on how much each action disrupts a specific adversary's strategy and how much damage it prevents if that disruption fails. The methodology is opinionated, quantitative and deliberately forces trade-offs into the open.

1. Introduction

Security teams face a persistent asymmetry. Threat intelligence continuously improves visibility into what adversaries do, while defensive planning must answer a fundamentally different question: what failures are unacceptable and which controls

meaningfully interfere with a specific adversary's strategy? These two questions share vocabulary but not structure.

In environments where time, budget or certainty are constrained, it is rarely feasible to implement every recommended control simultaneously. Prioritization is not a luxury but a structural constraint. Yet most prioritization today happens implicitly, inherited from compliance frameworks, risk heat maps or executive preference. When attacker intent shifts, these prioritizations rarely shift with it.

> *If defensive priorities do not change when attacker intent changes, prioritization is no longer strategic.*

ACDP was developed in response to this gap. It does not attempt to replace technique catalogs or maturity models. It provides a structured method for ordering defensive actions against a specific adversary profile, producing a prioritization that is explicit, defensible and adaptable when the threat landscape moves.

2. Positioning Relative to Existing Frameworks

Existing security frameworks fall broadly into three categories, each answering a different question.

2.1 Behavioral Catalogs

MITRE ATT&CK, D3FEND and similar frameworks describe what adversaries do at the technique level and what defensive measures exist to address those techniques. They are invaluable for behavioral modeling, detection engineering and post-incident analysis. By design, they describe coverage space rather than priority.

2.2 Control and Maturity Frameworks

NIST CSF, CIS Controls, ISO 27001 and equivalent compliance-oriented frameworks provide structured inventories of defensive measures, often with maturity levels. They address the question of what good practice looks like in aggregate. They do not answer which of those measures matters most against a particular adversary at a particular moment.

2.3 Risk Quantification Methods

FAIR and related quantitative risk methodologies translate risk into financial terms. They are rigorous on the impact side but agnostic about attacker intent. A control may reduce expected annualized loss without meaningfully interfering with a specific adversary's objectives.

ACDP operates one level above these frameworks. It consumes their outputs, technique catalogs, control inventories, impact assessments and produces an ordering that is actor-specific. It is not a replacement for any of them. It is the missing decision layer that translates threat intelligence and control catalogs into prioritized action.

3. The ACDP Perspective

The defining shift in ACDP is not methodological but perspectival. Three changes in framing distinguish it from conventional prioritization.

3.1 From Technique Coverage to Strategic Disruption

Conventional prioritization maximizes the number of adversary techniques a defender can detect or prevent. ACDP asks which controls most disrupt a specific adversary's strategy. A control that prevents one high-value technique may outrank a control that covers ten low-value techniques, depending on the adversary.

3.2 From Static Maturity to Adaptive Decision-Making

Maturity models imply a monotonic progression toward a single defensive ideal. ACDP rejects this assumption. When the threat landscape shifts, when a new adversary becomes relevant, when operational constraints change prioritization must shift with it. The output is not a roadmap but a decision at a point in time, reproducible as conditions change.

3.3 From Best Practice to Context-Driven Trade-Offs

Generic best practices assume idealized conditions. ACDP assumes real ones: finite budgets, scarce expertise, political constraints, legacy systems and adversaries who are actively choosing their tradecraft. It encourages thinking in terms of losing conditions rather than coverage metrics.

4. The Scoring Framework

ACDP evaluates each candidate control across four axes. Scoring is performed on a five-point ordinal scale, with explicit anchor descriptions to reduce inter-rater variance.

4.1 Axis A: Actor Disruption Value (ADV)

How strongly does this control interfere with the actor's campaign strategy? A maximum score indicates the control breaks or invalidates the adversary's chosen approach entirely. A minimum score indicates negligible effect on operations.

4.2 Axis B: Impact Risk Reduction (IRR)

If disruption fails and the adversary succeeds in reaching their objective, how much damage does this control prevent? A maximum score indicates prevention of catastrophic or irreversible outcomes. A minimum score indicates cosmetic reduction only.

4.3 Axis C: Cost and Operational Complexity (CC)

How realistic is implementation under current resource constraints? This axis captures financial cost, staffing requirements, operational disruption during rollout and ongoing maintenance burden. Higher scores indicate lower cost and lower friction.

4.4 Axis D: Detection-to-Decision Time (DDT)

Does the control provide usable signal early enough to change outcomes? A maximum score indicates pre-impact detection or prevention. A minimum score indicates the control provides only post-incident forensic value.

The four axes are deliberately chosen to cover the two dimensions that matter against any adversary, disruption of their plan and limitation of their damage plus the two dimensions that matter for implementation viability.

5. Actor-Specific Weighting

Axes are weighted according to the adversary being prioritized against. The weighting encodes strategic assumptions about the adversary's objectives and operational mode. For a destructive, state-aligned actor prioritizing strategic impact over access longevity, the recommended weighting is:

Axis	Weight	Rationale
ADV - Actor Disruption	**35%**	Disrupting strategy is primary goal
IRR - Impact Reduction	**35%**	Catastrophic outcomes must be bounded
CC - Cost & Complexity	**15%**	Feasibility constraint, not primary driver
DDT - Detection Time	**15%**	Detection valuable but recovery trumps it

The weighting is not fixed. Against a financially motivated crimeware operator prioritizing access persistence, DDT weighting rises significantly; against an espionage actor prioritizing stealth, IRR weighting drops relative to ADV. The key property of ACDP is not any specific weighting but the requirement that weightings be made explicit and justified against an adversary profile.

6. Worked Example: Sandworm

To illustrate how ACDP produces concrete priorities, this section walks through a scoring exercise against Sandworm, modeled as a destructive, state-aligned actor prioritizing strategic impact over access longevity.

6.1 Candidate Controls

Six controls were selected as representative of commonly recommended defensive investments:

- Immutable offline backups
- VPN and edge device inventory with patch SLAs
- Historical DNS and infrastructure analysis
- PowerShell command-line logging
- Scheduled task auditing
- Security awareness training

6.2 Raw Scores

Each control was scored across the four axes on the 1–5 scale:

Control	ADV	IRR	CC	DDT
Immutable offline backups	5	5	3	5
Edge inventory & patching	4	4	4	4
Historical DNS analysis	4	3	4	4
PowerShell logging	3	2	4	3
Scheduled task auditing	3	2	3	3
Awareness training	1	1	5	1

6.3 Priority Index Calculation

The Priority Index (PI) is calculated by applying the weighting vector for a destructive actor:

$$PI = (ADV \times 0.35) + (IRR \times 0.35) + (CC \times 0.15) + (DDT \times 0.15)$$

Applying this formula to each control yields the final prioritization:

Control	Priority Index	Tier
Immutable offline backups	**4.70**	Tier 1 - Immediate
Edge inventory & patching	**4.00**	Tier 1 - Immediate
Historical DNS analysis	**3.75**	Tier 2 - Planned
PowerShell logging	**2.80**	Tier 2 - Planned
Scheduled task auditing	**2.75**	Tier 3 - Deferred
Awareness training	**1.90**	Tier 4 - Deprioritized

7. Interpretation

The scoring produces several outcomes that conflict with conventional prioritization wisdom and this conflict is the primary analytical value of the methodology.

7.1 Recovery Outranks Detection

Immutable offline backups score highest because against a destructive actor, recovery capability is the only control that addresses the defining adversary objective. Detection controls that would be prioritized against espionage actors rank lower because they do not prevent the adversary's goal, only observe it.

7.2 Infrastructure Hygiene Outranks User Behavior

Edge inventory and patching outranks awareness training by a wide margin. Sandworm reliably gains initial access through unpatched internet-facing systems rather than through user interaction. For this actor, heavy investment in awareness training is a waste of resources regardless of how useful it might be elsewhere.

7.3 Boring Controls Dominate

The top-tier controls are not novel or visible. They are the kinds of investments that do not appear in executive presentations and do not generate vendor interest. ACDP forces recognition that defensive effectiveness and defensive visibility are weakly correlated.

7.4 Quantification Exposes Rationalization

Awareness training scores 1.90 against this adversary. Many organizations allocate substantial budget to awareness training while underinvesting in backups. ACDP does

not forbid this allocation; it requires that the allocation be defended in terms other than adversary-disruption value.

> *If you cannot explain your priorities numerically, you are not prioritizing. You are rationalizing.*

This claim requires qualification. ACDP does not make prioritization objective; it makes it explicit. A score of 5 rather than 4 remains a judgment. The value of the methodology is not that it removes judgment but that it documents judgment in a form that can be reviewed, contested and updated as evidence changes.

8. Limitations and Trade-Offs

ACDP is deliberately opinionated and carries corresponding limitations that practitioners should account for.

8.1 Dependency on Threat Intelligence Quality

Actor-centric prioritization is only as reliable as the threat intelligence that defines the actor profile. Incorrect assumptions about adversary intent produce systematically misaligned priorities. ACDP outputs should carry an explicit reference to the intelligence basis on which they were produced.

8.2 No Control Design Guidance

ACDP orders controls but does not design them. It assumes a candidate set of controls exists and that foundational security hygiene is in place. Organizations without baseline controls will find that ACDP presupposes work that still needs to be done.

8.3 Resistance to Exhaustive Coverage

The methodology forces decisions about what not to prioritize. Organizations accustomed to comprehensive compliance narratives may find this uncomfortable. Deprioritizing awareness training against Sandworm is defensible in the ACDP framework, but may be politically difficult in an organization that has invested heavily in such training.

8.4 Weighting as a Judgment Layer

The weighting vector encodes strategic assumptions that cannot be derived from data alone. Different analysts will produce different weightings for the same adversary profile. ACDP requires that these weightings be explicit and documented, but it does not and cannot eliminate the underlying judgment.

9. Integration with Broader Intelligence Workflows

ACDP is designed to be composable with actor-centric threat intelligence pipelines. When actor profiles, campaign observations and infrastructure modeling exist in a structured form for example, through continuous infrastructure tracking using an actor-centric intelligence graph, the inputs to ACDP scoring become both richer and more reproducible.

In the Malwarebox project, ACDP operates as the defensive translation layer above Kraken, the continuous observation platform and IIM, the Infrastructure Intelligence Model. Kraken provides the observational substrate, IIM provides structural grammar for describing how adversaries operate and ACDP translates the resulting intelligence into prioritized defensive decisions. The three components form an analytical loop: observation produces structure, structure produces priority and priority drives further observation.

ACDP does not require this integration to function. It can be applied with threat intelligence from any source. Integration amplifies its precision by making the inputs auditable and the scoring reproducible across teams and over time.

10. Conclusion

Security failures are rarely caused by missing controls. They are caused by misaligned priorities, by defensive investments that reflect compliance narratives, vendor influence or institutional inertia rather than the specific adversaries an organization faces.

ACDP addresses this misalignment by shifting focus from technique coverage to adversary intent and from abstract risk to tangible impact. Its value lies in deliberate selection. Used correctly, it does not replace existing frameworks but sharpens them by forcing uncomfortable but necessary decisions into the open.

The methodology is intentionally opinionated. It assumes explicit, challengeable priorities are better than implicit ones no one questions. Whether a given ACDP output is correct is a matter for review and debate. That the output exists in a form that permits review is the contribution of the method.

Feedback and critique on this methodology are welcomed. The author can be reached via the contact information on blog.synapticsystems.de.