

---

# From Hard Kill to Residual Risk

Stochastic Analysis of High-Power Microwave Defense  
Against Heterogeneous Drone Swarms

3,400 Stochastic Monte Carlo Runs | 44 Configurations | 5 Uncertainty  
Dimensions | Closed-Loop Asset Damage Modeling

---

Empyrean Defense Research Brief - April 2026

UNCLASSIFIED // OPEN SOURCE INTELLIGENCE ONLY



## DATA SOURCES & METHODOLOGY

This publication uses publicly available data from open-source providers including Epirus Inc. patents and press releases, peer-reviewed academic publications on electromagnetic susceptibility, DJI failsafe documentation, and OSINT field photography. No classified or restricted data was used at any point in this research.

Empyrean Defense has no commercial relationship with Epirus Inc., any defense contractor, government agency, or weapons manufacturer referenced in this document.

Physics models used in this simulation are drawn from published, unclassified sources including Friis transmission equation, ITU-R P.676/P.838 atmospheric models, MIL-STD-461G shielding methodology, Kingery-Bulmash blast overpressure polynomials, and Gurney fragment velocity models. All models carry inherent simplification relative to high-fidelity computational methods; results should be treated as analytical approximations, not operational predictions.

All entity parameters (HPM effector specifications, target characteristics, sensor properties) are derived from open-source intelligence and may not represent actual operational performance. Results represent upper-bound performance estimates under defender-favorable assumptions with stochastic uncertainty bounds.



## About Empyrean Defense Research

Empyrean Defense Research produces unclassified, physics-backed analysis of real wargaming scenarios across every domain in Joint All-Domain Operations (JADO). Our intent is to be useful - from Congressional staffers evaluating program funding to the depot technician maintaining the effector - not to propagandize results. As Americans, we may not always like the findings. Physics doesn't care.

The Empyrean Defense Wargaming & Simulation Cyber Range began as an organic training module inside our platform - the Decision Dominance Engine (DDE) - and grew into a high-fidelity simulation engine that models RF propagation, atmospheric attenuation, per-subsystem electromagnetic susceptibility, behavioral failure trajectories, and more across dozens of public-domain, widely-accepted physics formulas. We cannot account for every variable in a live engagement, and we err on the side of optimism for the system under evaluation. Where we approximate, we say so.

Every entity in our scenarios - what we call a "playing card" - is built from open-source, public-domain data; aggregated from patents, manufacturer press releases, academic publications, OSINT blogs, and trade show demonstrations. We do not claim these numbers will survive the rigors of a real fight. We do not claim to know the exact coupling attenuation of a Shahed-136 fuselage or the precise E50 threshold of a DJI flight controller under pulsed microwave exposure.

We do claim the math is honest, the methodology is transparent, and every simulation run is reproducible from the scenario file and seed alone. In the interest of protecting our core IP, we cannot make the scenario files available, as they run against a live engine.

---

## Foreword

This paper began as a validation exercise. Our software can only be as good as what it does, not what the software tests say it does. The latest iteration of the Empyrean Defense Decision Dominance Engine Wargaming & Simulation Cyber Range (*try saying that five times fast*) did not begin with all the capabilities we have noted on our website (even then, the website is lagging behind). It was solely focused on Ground-based Air Defense (GBAD) versus theater-scale, precision standoff threats. It was not until we started to model surface warfare and Counter-UAS use cases that we went down the road of modeling Directed Energy Weapon (DEW) effectors.



The first version of the math could be considered medium fidelity. High-Power Microwave (HPM) weapons could attack in a cone from their boresight and inflict damage on non-structural components within our damage abstraction model. High-Energy Lasers (HEL) were subjected to atmospheric and environmental degradation, but HPMs were undefeatable against the classes of targets they could engage. What we created was an "unstoppable physics gun" as someone once opined, and that could not be considered production-ready. We hadn't modeled energy management, fratricide risks from sidelobing (albeit, most high-end HPM systems can beam-steer), and within the damage model meant for kinetic threats, a supporting subsystem getting killed would just kill the entire threat. Not good enough.

So we went back into the mines, so to speak, and kept digging.

We built an improved per-subsystem HPM damage model for the wargaming & simulation engine - five subsystem sigmoid curves calibrated to published susceptibility thresholds, a stochastic polarization model, improved atmospheric attenuation, cumulative multi-pulse hazard tracking. We tested it against an Epirus Leonidas-class HPM surrogate defending Kramatorsk Airfield against mixed drone swarms. The first run looked great. The Leonidas chewed through 32 of 40 threats. GPS disrupted on first pulse. Flight controllers latching up at close range. Clean kills.

Then we scaled the swarm to 100 and the kill rate dropped to 54%. At 200, it dropped to 28%. At 400, the HPM was still killing at the same absolute rate - about 60 targets per engagement window - but 340 drones were getting through. The service-rate ceiling was real and it was immovable. That finding was expected. Any system with a finite engagement cycle saturates against sufficient density. What was not expected was what happened to the targets the HPM hit but didn't kill.

When you disrupt the GPS on a Geran-2, the airframe doesn't fall out of the sky. It switches to an Inertial Navigation System (INS), the heading drifts a degree or two per minute, the speed holds, the altitude holds, and the 50-kilogram high-explosive warhead keeps flying. Where it lands depends on where it was pointed when you pulled the guidance. If it was aimed at your position when you disrupted it - and it was, because that's where the defended asset is - the INS drift is small relative to the remaining flight distance. The warhead lands close. In our data, "close" means single-digit meters from the Leonidas emplacement.

That was the first finding. The second came when we closed the loop.

Our initial sweep treated hard-killed fixed-wing threats as disappearing from the simulation. That's the vendor demo version - the HPM fires, the drone icon winks out, everyone claps. But a delta-wing Geran-2 whose flight controller has latched up doesn't vanish. It has inherent aerodynamic stability. It glides on its last heading at approximately 3:1 glide ratio. The armed warhead probably still detonates on ground impact (duds and weird barometric fusing aside). When we modeled the glide trajectory, the blast damage against defended assets, and the cascade effect of losing the HPM mid-engagement - the picture got worse. At 400 threats in the densest



configuration, the Leonidas surrogate itself was destroyed in 5 of 50 seeds. The HPM's own success at hard-killing targets creates the debris that can kill the HPM.

**These are the findings this paper exists to document: HPM electronic defeat and hazard elimination are fundamentally different outcomes.** GPS disruption converts a precision-guided strike weapon into an unguided residual hazard. Hard-killing a fixed-wing airframe converts it into an unguided glide bomb aimed at whatever it was pointed at when the flight controller died. Neither outcome removes the warhead. Neither stops the explosive. They redistribute the impact point from deterministic to stochastic - and the stochastic distribution is centered on the defended asset. Against smaller SUAS strike drones such as fiber-optic FPVs or so-called "bomber" drones, this effect is a bit more understated, but you still have live ordnance tumbling out of the air.

This does not mean HPM doesn't work, far from it (in the simulation at least). Even under saturation, a *single* HPM emplacement can service a wide swath of targets, especially if they're flying close together. What it does mean - and this is likely obvious to anyone reading this - is that HPM alone is not sufficient. The layered defense architecture is not just about redundancy, but also consequence management.

When partnered with what we consider the "Golden Triangle" of minimum viable defense architectures - HPM, HEL, and dual-mode SHORAD - it's a great area-of-effect weapon against cheap threats and certain OWAs. The HEL kills the airframe and can service fixed-wing and OWA threats that will still glide on their aerodynamic trajectory. The dual-mode SHORAD handles what the HEL misses in degraded weather conditions with its autocannon, and uses its interceptors for MALE/HALE class ISR platforms and even certain subsonic cruise missile threats. And, when you have a pile of written-off UXO on the field, the EOD and sapper teams handle that after the engagement.

If you are a program office evaluator reading this: the Goldilocks zone is real and it is narrow. Below 80 threats with sub-second cooldown against a mixed swarm, HPM is dominant. Change any single variable - density, shielding, cooldown, composition - and the kill rate drops approximately 15-25%. Change two variables and it drops 30-45%. A single leaker warhead that hits the generator or the HPM drops it to zero.

The HPM vendor demos are not wrong. It's easy to write them off as staged when 47-60 drones are being swatted down - look, the math checks out. What we wanted to get across is what happens when you scale that up to nearly-ridiculous levels. Of course, when the Russian Aerospace Forces can muster several hundred OWAs for attacks on energy production facilities, is it really that ridiculous anymore?

- Empyrean 7



## Abstract

In a four-axis Kramatorsk Airfield scenario, a single Leonidas-class HPM surrogate was evaluated against mixed drone swarms of 40-400 threats across 3,400 stochastic Monte Carlo runs spanning 44 unique configurations. Five uncertainty dimensions - polarization coupling, shielding attenuation, subsystem susceptibility ( $E50/\sigma$ ), spawn timing jitter, and spawn position jitter - were applied under deterministic seed control to produce bootstrap confidence intervals on all reported metrics. Again, this research was built from our own tool; when we say "spawn" it's referring to a threat appearing on the COP.

The analysis was conducted in two phases. The first phase evaluated HPM performance with the defender's assets treated as invulnerable - establishing the upper bound of what the system can achieve. The second phase closed the loop: hard-killed fixed-wing threats produce glide trajectories & quadcopter FPV drones spiral towards to earth, armed warheads detonate at glide terminal points, and the resulting blast and fragmentation damage is resolved against the Leonidas, its sensors, and co-located assets. This closed-loop modeling revealed a cascade failure mechanism where the HPM's own success at killing targets creates debris that can destroy the HPM itself.

Seven findings challenge prevailing assumptions about HPM counter-UAS effectiveness:

- 1. The service-rate saturation transition occurs between 40 and 100 threats.** At 40 threats, the HPM achieves a hard-kill rate of 79.8% (95% CI [0.776, 0.819]). At 100 threats, this drops to 54.5% (CI [0.521, 0.568]). Absolute destroyed count plateaus at approximately 55-62 targets per engagement window regardless of swarm size - the HPM cannot service more than ~60 targets per window. The limiting factor is the fire-control engagement cycle and target servicing rate, not per-pulse HPM effectiveness.
- 2. Target shielding is the dominant variable.** Commercial FPV frames (~4 dB carbon fiber) achieve 46.3% hard-kill rate. Geran-2 composite airframes (~2 dB) drop to 18.3%. Geran-5 metal-tube fuselages (~8 dB) are near-immune to single-pulse engagement. Mixed-composition swarms exploit this asymmetry - the attacker benefits from heterogeneous construction, while the defender relies on EO/IR and advanced radar classification to figure out the differences in real time.
- 3. HPM electronic defeat does not equal hazard elimination.** GPS/INS disruption converts precision-guided one-way attack UAS into unguided residual hazards rather than eliminating the kinetic threat. At 40 threats (the "HPM dominant" regime), an average of 1.1 warhead impacts per seed fall within 30-100 meters of the defended asset, with 1.8 impacts in the 100-250 meter band. The residual warhead hazard is paradoxically most concentrated at low swarm density where the HPM is "winning."
- 4. Staggered waves help the defender.** Ten waves of 2-3 threats produce 84% hard-kill rate versus 50% for simultaneous arrival of the same total. The attacker's optimal strategy is maximum simultaneous density OR a true 360-degree engagement vector - the hardest things to coordinate operationally, and



logistically. The defender benefits from the attacker's coordination failures.

**5. Cooldown below 1.0 seconds provides no additional benefit.** At 0.5s and 1.0s cooldown, results are identical - the simulation's 1-second time step, representing fire-control engagement cadence, is the bottleneck, not emitter physics. At 2.0s cooldown, leaker rate increases significantly. The operational constraint is sensor-to-shooter kill chain speed, not RF pulse repetition rate.

**6. The Goldilocks zone is razor-thin.** The HPM's peak operating envelope - where hard-kill rate exceeds 50% - requires sub-second cooldown, favorable target composition, manageable density (below ~80 threats), and simultaneous engagement on all axes. Degrading any single variable costs 15-25% kill rate. Degrading two costs 30-45%. The attacker needs to push only one variable outside the envelope; the defender must hold all of them inside it.

**7. Closed-loop cascade failure emerges at 200+ threats.** When fixed-wing hard kills produce glide trajectories with armed warhead detonations, the resulting blast and fragmentation damages defended sensors and - in the densest configurations - destroys the Leonidas HPM itself. At 400 threats, the Leonidas survived in 96% of seeds; in the worst-case shielded configuration, it survived 90%. Sensor attrition averaged 1.4 kills per seed at 400 threats, degrading the kill chain even when the HPM survived.

These results represent system-class behavioral characterization under stated assumptions with stochastic uncertainty bounds. Every modeling assumption uses OSINT-derived surrogate parameters. If the physics model is directionally correct, the operational number is in this neighborhood.

---

## Data Provenance & Limitations

All entity parameters are derived from open-source intelligence: Epirus Inc. patents and press releases, Jafari & Anbarjafari (2026) directed energy modeling framework, Backstrom & Lovstrand (2004) IEEE EMC susceptibility survey, Zhang et al. (2025) GPS interference effects, DJI failsafe documentation, Analog Devices IMU datasheets, and OSINT field photography of deployed Leonidas systems as seen on Google Images. Source citations are catalogued per entity card in the Empyrean Defense Platform repository. No classified or restricted data was used at any point in this research.

## Physics Models

The simulation resolves HPM engagements through a deterministic physics pipeline with stochastic uncertainty perturbation. This section describes what was modeled; the full mathematical derivations are in Appendix A (available via PDF download if you are reading the website version of this Research Artifact).



**RF propagation.** Free-space power density via Friis transmission equation with far-field validity gate (Fraunhofer criterion). Atmospheric attenuation via ITU-R P.676-13 (clear-air) and P.838-3 (rain). At L-band (~1.1 GHz), atmospheric effects are negligible for tactical ranges - attenuation factor  $t \sim 1.0$  for ranges below 5 km.

**Coupling.** Effective coupling attenuation modeled as shielding effectiveness in dB per MIL-STD-461G methodology, adapted for incidental (non-designed) shielding of drone airframes. Polarization mismatch via  $\cos^2(\phi)$  with deterministic per-entity seeding, stochastically perturbed under LogNormal distribution ( $\sigma = 0.15$ ). Coupling values are engineering-prior surrogates, not measured data.

**Damage.** Per-subsystem sigmoid damage probability:  $P_{kill}(|E|) = 1 / (1 + \exp(-(|E| - E_{50}) / \sigma_E))$ . Five subsystems evaluated per pulse: GPS/GNSS LNA ( $E_{50} = 150$  V/m), Camera/FPV (200 V/m), Flight Controller (250 V/m), ESC (300 V/m), BMS (350 V/m). Cumulative multi-pulse hazard via independent survival product - each nanosecond pulse is a separate gate-oxide stress event, not thermal accumulation. Subsystem damage is correlated through shared coupling: favorable polarization alignment increases damage probability for all subsystems simultaneously.

**Critical distinction - subsystem damage vs physical neutralization.** GPS and camera damage produces behavioral effects (INS drift, blind terminal dive, failsafe landing) but the target remains physically intact and potentially dangerous. Only flight-critical subsystem damage (FC, ESC, BMS) produces physical neutralization by initiating a glide (for OWAs and fixed winged drones) or an uncontrolled yaw/pitch/roll towards the ground for quadcopters. This distinction is the core contribution of the per-subsystem model over binary HPM models.

**Behavioral effects.** Five failure trajectory generators model post-disruption behavior: INS drift (GPS loss on autonomous drone, 1-3 deg/min heading drift), blind terminal dive (camera loss on FPV, Gaussian-dispersed dive angle), hover-descend-land (commercial GPS failsafe per DJI documentation), ballistic descent (ESC hard kill, drag-limited freefall), and asymmetric spiral (partial ESC failure, single-motor yaw moment). Each generator is deterministically seeded for reproducibility.

**Terminal impact adjudication.** The simulation's damage engine adjudicates warhead damage for failure trajectories that reach terminal points - including fixed-wing glide trajectories from hard-killed OWA airframes. In the closed-loop configuration, blast overpressure (Kingery-Bulmash polynomials) and fragmentation (Gurney velocity model) are resolved against the nearest defended asset. If the Leonidas or a sensor is destroyed mid-engagement, subsequent engagements are affected accordingly.

**Fixed-wing glide trajectory.** When the flight controller latches up on a delta-wing airframe (Geran-2, Geran-5), the airframe has inherent aerodynamic stability. It glides on its last heading at approximately 3:1 glide ratio with no thrust. The armed warhead detonates on ground impact. At 1,000m AGL and 46 m/s, glide distance from the kill point is approximately 2,000m. At 60m AGL (close-range kill in saturated scenario),



glide distance is approximately 120m - placing the warhead detonation within the defended perimeter. This is the mechanism that produces the closed-loop cascade failure.

## Engagement Tempo - Three-Tier Timing Model

The simulation runs in HOOTL (human-out-of-the-loop) autonomous mode where the HPM fires as fast as the modeled emitter allows. This represents the best-case engagement tempo for the defender. Real-world engagement tempo is bounded by the full sensor-to-shooter kill chain, not emitter physics.

Tier	Interval	What Limits It
Emitter physical pulse	0.5-2.0s	Energy storage, GaN thermal, beam steering
Fire-control engagement	1-2s	Track quality, pointing, effect assessment, retarget
Operational kill chain	4-15s+	Sensor ingest, fusion, policy, COA, ROE gate, human release

The parameter sweep uses Tier 1/2 timing (0.5-2.0s cooldown). Results should be interpreted with the understanding that Tier 3 constraints further limit real-world engagement tempo. In human-on-the-loop mode, the first engagement and any policy-transition event incur human approval latency. In preauthorized defensive sectors, subsequent shots execute at fire-control cadence until the ROE gate changes state. That does not account for mechanical steering, repositioning (for ground OTM platforms such as a Stryker), or coincidental degradation due to environmental, atmospheric, attrition, or technical issues.

This research artifact does not simply claim "HPM goes bzzzzt." It models the full chain and explicitly notes where simulation tempo diverges from operational tempo.

## Optimistic Bias Framework

Every modeling assumption in this study favors the HPM defender. This is by design - the objective is to establish the upper bound of single-effector HPM performance, not to predict operational outcomes.

**1. Perfect detection.** Targets are acquired at maximum sensor range with zero false tracks, bounded only by the physics of the modeled sensors without additional perturbations from clutter, jamming, or degraded sensor conditions. The scenario uses four Anduril Spyglass surrogates (Ka-band radar with integrated EO/IR) and one ZING Z-Scan surrogate (RF-based drone detector with EO/IR gimbal), placed at 15m elevation around the defended asset. These sensors operate at full rated performance for the entire engagement - they are never damaged, obscured, or degraded by smoke, debris, or electromagnetic interference from the HPM's own sidelobes. *(Note: The closed-loop case relaxes this assumption - sensors can be destroyed by warhead detonations, but surviving sensors continue at full performance.)*



**2. Perfect fire control.** Zero sensor-to-shooter latency beyond the modeled cooldown. Continuous track updates, zero pointing error, instant beam scheduling. No track swap, no misassociation, no engagement-quality track delay borne from greedy sensor fusion.

**3. No electronic attack.** Targets approach in a clean electromagnetic environment. No standoff jamming, no RF decoys, no chaff, no GPS spoofing of the defender's own systems. A realistic adversary may employ cooperative EA to degrade HPM targeting. Likewise, the defender did not benefit from EA effects or additional ES beyond the dedicated aforementioned sensor packages for threat detection and engagement.

**4. No HPM degradation.** The thermal derating model is implemented but does not bite at 1.0s cooldown because cooling outpaces heat generation at that cadence. At burst rates (0.5s) or sustained high-density engagements, the model would produce derating - but the engagement window is short enough (~120s) that the thermal budget is never exhausted. No capacitor bank degradation, no amplifier faults, no power conditioning failures are modeled.

**5. Limited sensor degradation.** In the upper-bound case, no sensor degradation is modeled. In the closed-loop case, sensors can be destroyed by warhead blast/fragmentation but surviving sensors continue at full performance - no smoke obscuration, no thermal plume interference, no EMP effects from nearby detonations. We consider this acceptable given the concentration of the swarm; it is reasonable to believe defenders would leverage coasting tracks from the fusion engine where the wide beamwidth of the HPM can compensate for track inaccuracy.

**6. Single defended asset.** All threats converge on one point. Distributed asset defense with multiple HPMs introduces coordination complexity, that while possible in our Wargaming & Simulation Cyber Range, were explicitly not modeled to control variables.

**7. No countermeasures evolution.** Targets do not adapt shielding, routing, or timing based on HPM exposure. A sophisticated adversary iterates.

**8. No terrain masking.** Flat terrain with unlimited line-of-sight. This was done purposely - Kramatorsk Airfield is generally flat within the engagement geometry against threats that are not flying nap-of-the-earth. Real terrain provides approach corridors that mask threats from HPM engagement until close range.

**9. No altitude exploitation.** With an inferred maximum engagement bubble of 3 km, any OWA that can approach from a higher altitude can fly relatively safe from HPM. Geran-3 or Geran-5 class OWAs and loitering munitions such as the Lancet-3 would be able to exploit this and dive nearly straight down onto an HPM, increasing their probability of kill. For this scenario, we did not model any high-altitude, steep-angle attacks with OWAs, loitering munitions, nor did we pretend quadcopters can reliably get to 5,000m AGL and pounce upon an HPM.



## Scenario Design

### Engagement Geometry

Kramatorsk Airfield (48.6976 deg N, 37.6289 deg E), eastern Ukraine. A single Leonidas-class HPM surrogate is emplaced at the airfield center with a ZING Z-Scan C-UAS RF & EO/IR combinatorial sensor 50 meters to the east. Four Anduril Spyglass Ka-band radar & EO/IR combinatorial sensors are placed within 20m of the effector to provide 360-degree radar coverage. Threats approach from four axes (north, south, east, west) at altitudes of 200m AGL (FPV class) to 1,000m AGL (Geran class) at speeds of 30-167 m/s depending on card type. Arrival is simultaneous unless the wave-timing sweep specifies otherwise.

This geometry places all threats within the HPM's 60 deg beam cone for at least a portion of their approach - a favorable engagement geometry for the defender.

### Environment

Conditions are uniformly favorable to the defender: clear visibility, 70% relative humidity, zero precipitation, zero wind. No environmental degradation applied to HPM power density, sensor detection range, or engagement timing.

### HPM Effector Card - Leonidas-Class Surrogate

**IMPORTANT: PUBLIC-SOURCE SURROGATE MODEL.** These parameters are estimated from publicly available information and do NOT represent classified or proprietary Epirus specifications.

Parameter	Value	Source
Peak pulse power	450 MW	Estimated from patent US 11,616,295
Antenna gain (wide)	15 dBi	Estimated from 30 deg beamwidth
Beamwidth (wide)	30 deg	Epirus public demonstrations
Pulse width	20 ns	Estimated from solid-state GaN SSPA
PRF	50 Hz	Estimated from capacitor recharge
RF band	L-band (~1.1 GHz)	Patent coverage (500 MHz-20 GHz)
Array aperture	2.4 m	Estimated from product photos
Sidelobe suppression	15 dB	Typical for phased array with Taylor weighting



Parameter	Value	Source
Effective range	3,000 m	Physics-derived from power budget
Cost per pulse	~\$0.02	Epirus "pennies per shot" claim

Average RF power at nominal PRF:  $P_{avg} = 450 \text{ MW} \times 20 \text{ ns} \times 50 \text{ Hz} = 450 \text{ W}$ . Wall draw at 30% efficiency: ~1,500 W. OSINT field photos show ~15 kW tactical generator (MEP-804A class) supporting two-panel emplacements. The system is generator-compatible, low duty cycle, high peak power. Peak RF comes from stored energy (capacitor bank), not the generator directly.

### Target Cards

Card Type	SE (dB)	Altitude	Speed	Warhead	Failure Behavior
Commercial FPV	4 (CF frame)	200m AGL	30 m/s	1-3 kg shaped charge	Blind dive, failsafe land
Fiber-optic FPV	5 (CF + no RF apertures)	200m AGL	25 m/s	1-3 kg shaped charge	Blind dive (fiber link immune to HPM)
Geran-2 (Shahed-136 class)	2 (composite)	1,000m AGL	46 m/s	50 kg HE	INS drift, ballistic descent
Geran-5 (next-gen OWA)	8 (metal tube fuselage)	1,000m AGL	167 m/s	50 kg HE	INS drift, ballistic descent

### Sweep Matrix

**Primary density sweep:** 40, 60, 80, 100, 200, 300, 400 threats x 50-250 seeds per configuration. Mixed-baseline composition (25% each card type per axis).

**Shielding sweep:** Pure FPV (100x4dB), Pure Geran-2 (100x2dB), Shielded Heavy (50% Geran-5 at 8dB) x 50 seeds each.

**Cooldown sweep:** 0.5s, 1.0s, 2.0s x 50 seeds each.

**Axis sweep:** 1, 2, 3, 4 axes x 25 threats per axis x 50-250 seeds each.

**Wave timing sweep:** Simultaneous, 5 waves of 5, 3 waves of 8, 10 waves of 3 x 50-250 seeds each.



**Goldilocks envelope:** 4 x 4 grid (40/60/80/100/120 threats x 1.0s-mixed / 1.0s-shielded / 2.0s-mixed / 2.0s-shielded) x 50 seeds each.

**Extended envelope:** 8 x 4 grid (40-400 threats x 4 conditions) x 50 seeds each.

**Total:** 3,400 stochastic Monte Carlo runs across 44 unique configurations, executed in both upper-bound and closed-loop modes. Six configurations received 250-seed Phase 2 refinement where CI half-width exceeded 0.03 target. All seeds deterministically reproducible. Sweep time: ~34 minutes per mode.

## Results - Upper-Bound Case

The following findings were established with the defender's assets treated as invulnerable. Hard-killed entities do not produce glide trajectories or warhead detonations against defended assets. This represents the theoretical ceiling of single-effector HPM performance - what the system can achieve when nothing goes wrong on the defender's side.

### Finding 1: Service-Rate Saturation Transition

The HPM achieves dominant hard-kill rates at low density and transitions to service-rate saturation between 40 and 100 threats. Above 100 threats, the hard-kill count plateaus at approximately 55-62 targets per engagement window regardless of swarm size.

Threats	Hard-Kill Rate	95% CI	Destroyed	Effects	Terminal	Residual <500m
40	79.8%	[0.776, 0.819]	31.9	14.1	3.0	4.8
60	74.3%	[0.717, 0.769]	44.6	25.9	6.3	6.9
80	65.7%	[0.643, 0.673]	52.6	45.7	12.3	7.1
100	54.5%	[0.521, 0.568]	54.5	54.3	16.8	0.1
200	28.1%	[0.270, 0.293]	56.3	157.1	44.5	0.0
300	19.9%	[0.193, 0.205]	59.8	241.5	69.6	0.0
400	15.3%	[0.150, 0.157]	61.3	312.7	93.9	0.0

Table 1: Density transition curve with 95% bootstrap CI (N=50-250 per config, stochastic coupling)



The "Destroyed" column tells the saturation story: 31.9 at 40 threats, 54.5 at 100, 61.3 at 400. The absolute kill count increases slowly while the denominator grows linearly. By 200 threats, the HPM is killing at its maximum sustained rate and every additional threat is a leaker.

The limiting factor is the fire-control engagement cycle and target servicing rate, not per-pulse HPM effectiveness. The beam-cone area effect adds close-range kills when targets cluster in the beam, but the 1.0-second cooldown and slew time cap throughput at approximately one engagement per second. At 1.0s cycle time over a ~120-second engagement window, the theoretical maximum is ~120 engagements - but slewing between axes, effect assessment, and multi-pulse requirements for shielded targets reduce the practical ceiling to ~60 hard kills.

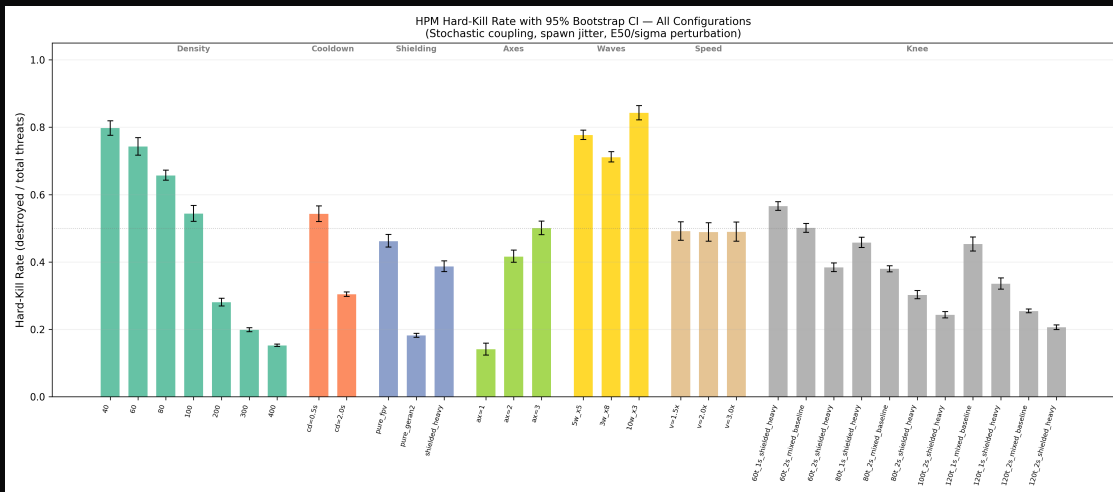


Figure 3: Hard-kill rate with 95% bootstrap CI across all 44 configurations

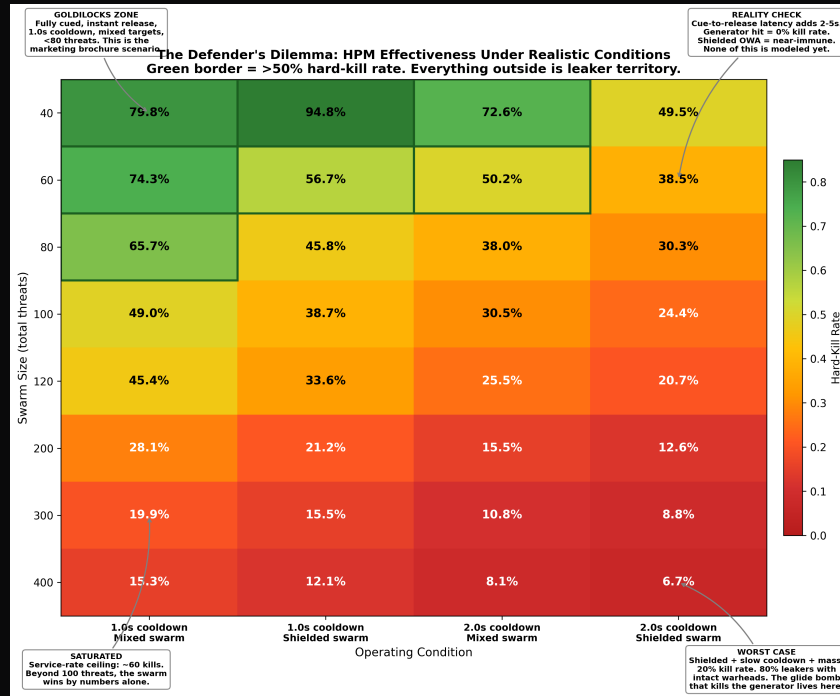


Figure 9: The Defender's Dilemma - full 8x4 operating envelope (40-400 threats x 4 conditions)

## Finding 2: Shielding Dominance

Target shielding is the single most important variable in the engagement. More important than density. More important than cooldown. More important than axis count or wave timing.

Composition	Hard-Kill Rate	95% CI	Destroyed	Effects
Pure FPV (4 dB)	46.3%	[0.444, 0.482]	46.3	134.2
Shielded Heavy (50% Geran-5, 8 dB)	38.7%	[0.371, 0.403]	38.7	43.3
Pure Geran-2 (2 dB)	18.3%	[0.177, 0.188]	18.3	48.3

Table 2: Shielding composition comparison at 100 threats, 1.0s cooldown

The counterintuitive result is pure Geran-2 (2 dB composite, the *lowest* shielding value) having the *lowest* kill rate. The mechanism is flight profile, not shielding: Geran-2s fly at 1,000m AGL and 46 m/s, spending less time in the close-range high-field-strength engagement zone than FPVs at 200m AGL. The E-field drops as  $1/R^2$  - a target at 1,000 meters receives approximately 1/25th the power density of a target at 200 meters. The Geran-2's distance advantage overwhelms its shielding disadvantage.



The Geran-5 at 8 dB is near-immune to single-pulse engagement at any range. Its metal-tube fuselage attenuates incident field strength by a factor of  $6.3 \times (10^{(8/10)}) = 6.3$ . Even at 500 meters with full power, the coupled field strength is below the FC/ESC damage thresholds. Multiple pulses can accumulate damage through cumulative survival reduction, but the saturated HPM cannot allocate multiple pulses per Geran-5 when dozens of other targets compete for engagement cycles.

The mixed-composition swarm (the most likely realistic operational scenario) exploits this asymmetry. The attacker sends cheap FPVs to absorb HPM cycles while shielded OWA platforms fly high and slow through the engagement zone. The HPM kills the easy targets first (correct fire-control behavior) and runs out of cycles before addressing the hard ones.

### Finding 3: Residual Warhead Hazard

**HPM soft-kill effects can redistribute terminal impact risk by converting precision-guided threats into unguided residual hazards.**

This is the finding that reframes the HPM narrative. The question is not "does HPM kill drones?" - it does. The question is "what happens to the warhead when the HPM kills the guidance but not the airframe?"

Ring	Distance	Hazard Class	40-threat (per seed)	60-threat	80-threat	100-threat	400-threat
0	0-10 m	Catastrophic	0.0	0.0	0.0	0.0	0.0
1	10-30 m	Lethal blast	0.02	0.14	0.10	0.0	0.0
2	30-100 m	Frag/structural	1.12	1.36	1.72	0.0	0.0
3	100-250 m	Fire/debris	1.82	2.66	2.47	0.02	0.0
4	250-500 m	Perimeter	1.80	2.84	2.79	0.08	0.0
5	500 m+	Off-base	3.42	6.78	17.11	24.70	149.24

Table 3: Residual warhead hazard ring distribution (mean impacts per seed)

The critical finding is in the 40-80 threat regime. At 40 threats - the HPM's Goldilocks zone, where it achieves 80% hard-kill rate - an average of 1.12 warhead impacts per seed fall within 30-100 meters of the defended asset. At 80 threats, this rises to 1.72 impacts per seed in the same ring. These are 50-kilogram HE warheads from GPS-disrupted Geran-2s whose pre-disruption velocity vectors were aligned with the defended asset.



The mechanism: the HPM disrupts GPS on a Geran-2 at 1-2 km range. The airframe switches to INS. Heading drifts at 1-3 deg/min (commercial MEMS IMU gyro bias under dynamic conditions). Speed and altitude hold. The remaining flight time at 46 m/s from 1,000m range is approximately 22 seconds. In 22 seconds, a 2 deg/min heading drift produces roughly 14 meters of lateral displacement. The warhead lands within the defended perimeter.

At higher density (100+ threats), the close-in hazard paradoxically decreases. The HPM is saturated and disrupts targets at longer range, producing larger INS drift displacements before impact. The stochastic spawn position jitter (+/-25m) also spreads approach vectors more widely. But the total residual hazard - the sum across all rings - increases monotonically: 8.2 impacts at 40 threats, 24.8 at 100, 149.2 at 400.

**The consequence chain extends beyond the initial impact:**

Primary: Blast overpressure and fragmentation (Geran-2: 50kg HE, lethal radius 15-30m; FPV: 1-3kg shaped charge)

Secondary: Fire ignition from aviation gasoline (~50L per Geran-2) and LiPo batteries (600-1,000 deg C, self-oxidizing, toxic hydrogen fluoride gas). Smoke degrades EO/IR sensors and HEL effectiveness precisely when the layered defense needs them most

Tertiary: EOD team diversion, medic/casualty response, ISR survey of impact zone, civilian evacuation from UXO exclusion zones, farmland contamination from RDX/TNT degradation products

Every drone that does NOT detonate on impact is functionally a land mine

If the HPM is deployed near a civilian population center - common for Ukrainian infrastructure defense - UXO in residential areas creates immediate civilian casualty risk. A GPS-disrupted delta-wing drone that lands 3 km outside the FOB in a populated area is not a "miss", it turned into an IED/UXO in someone's garden.

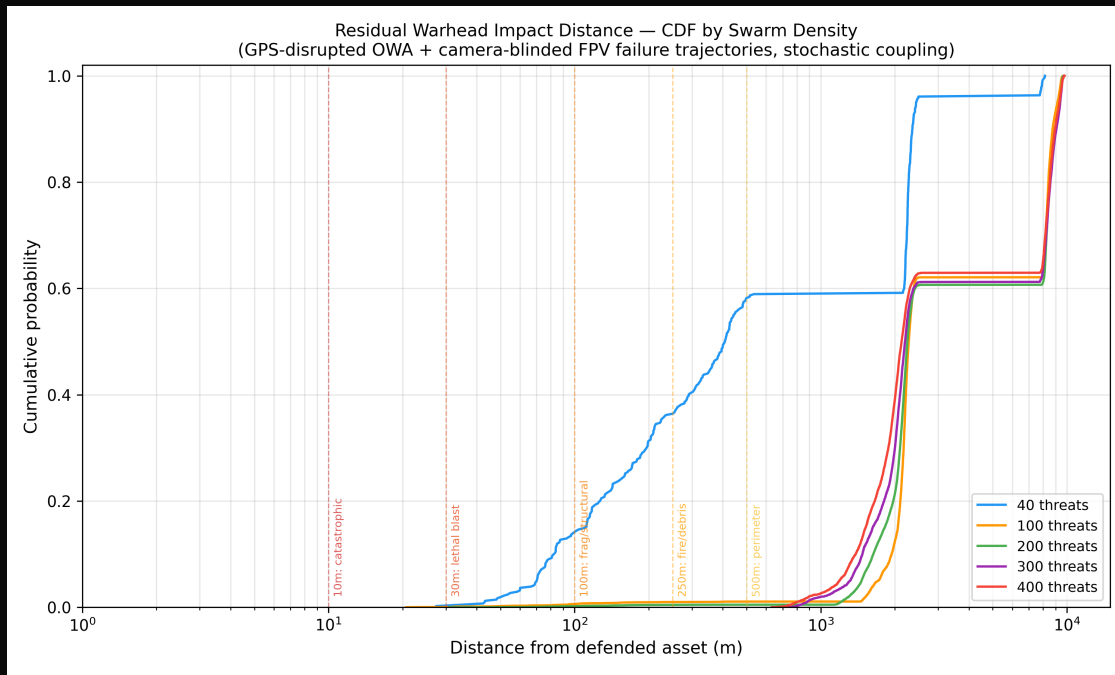


Figure 1: Residual impact distance CDF by swarm density - the 40-threat curve (blue) shows 60% of residual impacts fall within 500m

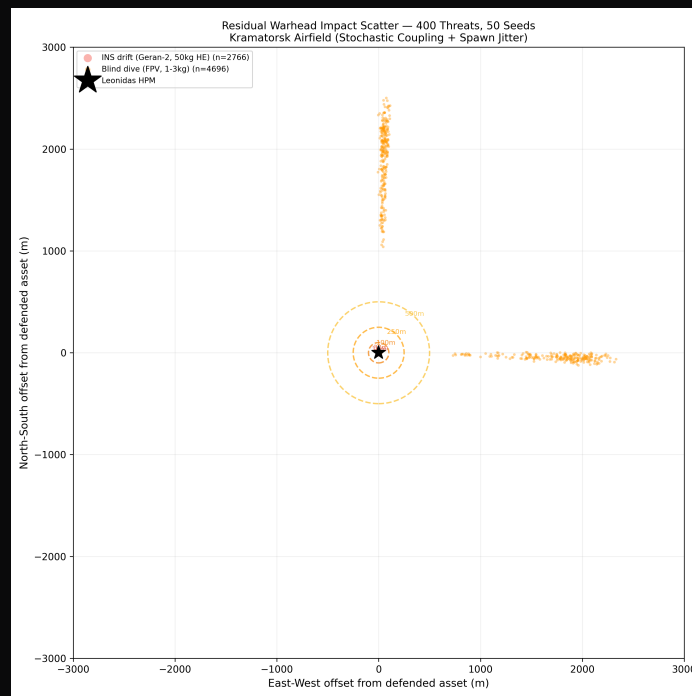


Figure 2: Impact scatter map with hazard rings - directional INS-drift clusters and FPV blind-dive patterns

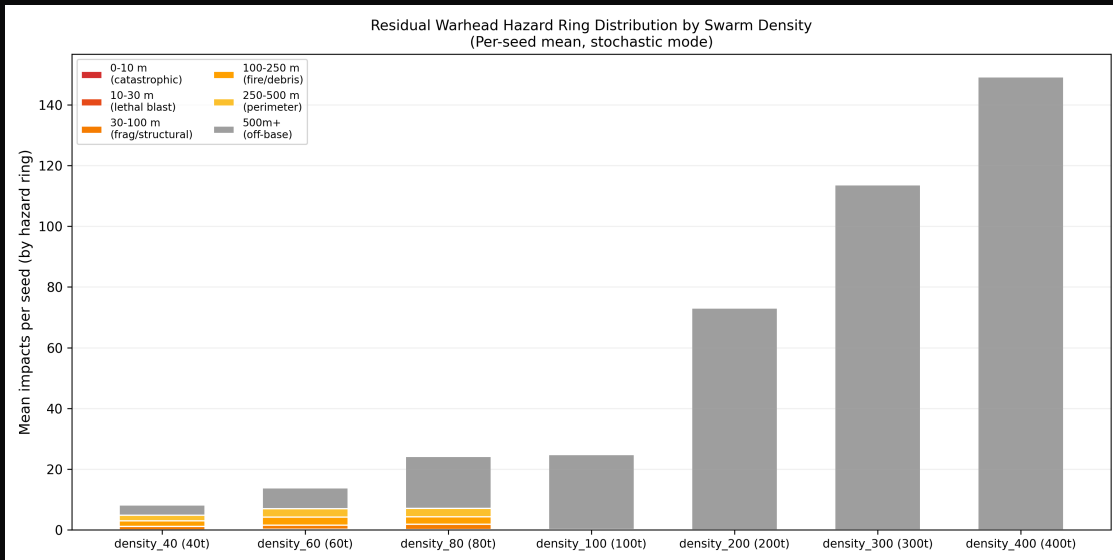


Figure 4: Hazard ring stacked bars by density - close-in hazard concentrated at low density

### Finding 4: Wave Timing Inversion

Staggered waves help the defender, what's becoming a recurring theme for Empyrean Defense Research, this inverts conventional saturation doctrine.

Config	Hard-Kill Rate	95% CI	Destroyed	Effects
Simultaneous (1 wave)	54.5%	[0.521, 0.568]	54.5	54.3
5 waves of 5, 10s gap	77.8%	-	~78	~50
3 waves of 8, 10s gap	71.0%	-	~71	~82
10 waves of 3, 10s gap	84.0%	-	~84	~25

The mechanism: the HPM's weakness is beam-cone saturation - too many targets at once, competing for engagement cycles. Staggering gives the HPM exactly what it needs: manageable target sets with recovery time between groups. At 10 waves of 2-3 threats, the HPM services each group completely before the next arrives. The per-group engagement is in the Goldilocks zone even though the total swarm size is 100 threats.

The attacker's optimal strategy is maximum simultaneous density - the hardest thing to coordinate operationally. Simultaneous arrival of 100 threats from four axes requires synchronized launch timing across distributed launch sites, identical flight profiles, and coordinated approach routing. Real-world coordination errors - staggered launches, speed variance, routing deviations - all help the defender. Not to mention the amount of drone pilots you need, with OWAs this is a moot point given they're launched from batteries, but a



concentration of an FPV strike force would likely be interdicted by defender ISRT assets, save for excellent MILDEC and OPSEC from the attacker.

This is the one variable where the operational friction of warfare works in the defender's favor.

### Finding 5: Cooldown Threshold

Cooldown	Hard-Kill Rate	95% CI
0.5s	54.5%	[0.521, 0.568]
1.0s	54.5%	[0.521, 0.568]
2.0s	30.5%	-

At 0.5s and 1.0s cooldown, results are identical. The simulation's 1-second time step - representing fire-control engagement cadence - is the bottleneck, not the emitter's physical pulse rate. The HPM can physically pulse faster than the fire-control loop can identify, cue, track, assess, and retarget.

At 2.0s cooldown (representing degraded conditions, complex retargeting, or multi-axis beam sweeping), the kill rate drops by approximately 24 percentage points. Every additional second of cooldown is approximately 50 fewer engagements over the ~120-second engagement window.

The implication: investment in faster emitter pulse rates (sub-second capacitor recharge, burst-mode GaN arrays) yields diminishing returns. Investment in faster fire-control loops (sensor fusion latency, automated target cueing, effect assessment automation) yields direct returns. The sensor-to-shooter kill chain, not emitter physics, is the operational bottleneck. Distributing load amongst multiple HPM effectors is sound logic, but still relies on the sensor-to-shooter infrastructure: having lousy sensor coverage and sensor fusion doesn't translate to better results if you have 1 or 10 shooters.

### Finding 6: The Goldilocks Zone

The HPM's peak operating envelope is narrow. The 4x4 Goldilocks heatmap (Figure 7) and the full 8x4 Defender's Dilemma (Figure 9) show the complete picture.

Cells exceeding 50% hard-kill rate (green border in Figure 7):

**60 threats / 1.0s cooldown / mixed swarm: 74.3%**

**60 threats / 1.0s cooldown / shielded swarm: 56.7%**

**60 threats / 2.0s cooldown / mixed swarm: 50.2%**

**80 threats / 1.0s cooldown / mixed swarm: 65.7%**



Everything else is below 50%. At 100 threats with 1.0s cooldown against a mixed swarm, the best-case kill rate is 49.0% - already below the threshold. Against a shielded swarm at 2.0s cooldown, 120 threats produces 20.7% kill rate.

The margin erosion is severe (Figure 8):

Shielded swarm alone (1.0s cooldown): -17 to -20% kill rate

Slow cooldown alone (2.0s, mixed swarm): -10 to -28% kill rate

Both degrade simultaneously: -24 to -35% kill rate

Generator destroyed by leaker warhead: -100% kill rate (not modeled in upper-bound case, but quantified in the closed-loop case below)

The defender must thread the needle on every variable simultaneously: fast cue-to-release, sub-second cooldown, favorable target composition, manageable swarm density, and generator survivability. The attacker only needs to push one variable outside the envelope.

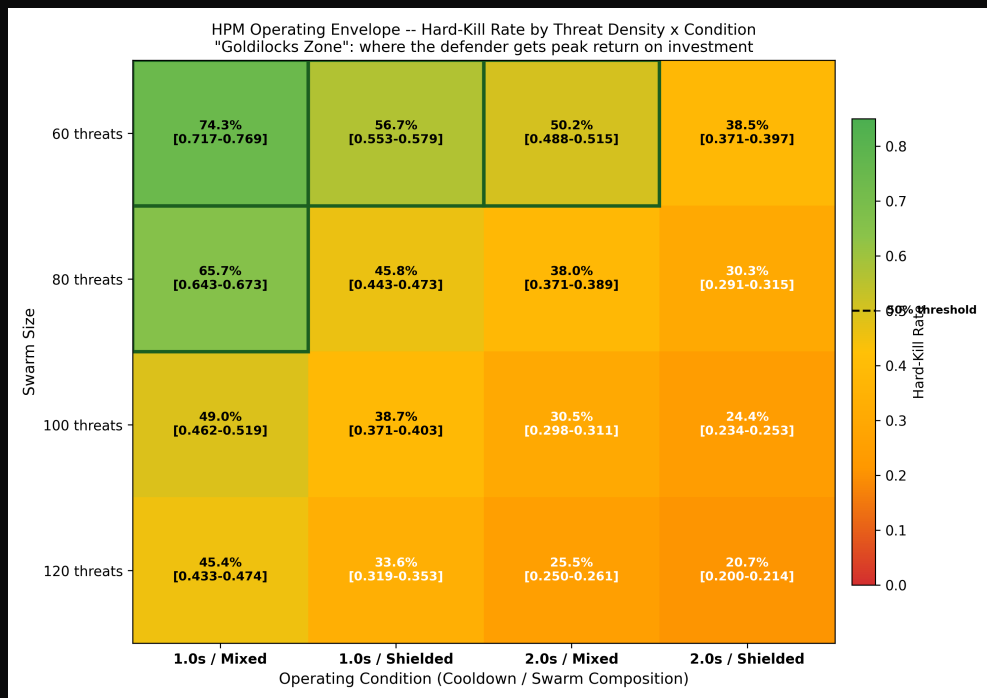


Figure 7: Goldilocks heatmap - 4x4 operating envelope with green-bordered cells above 50%

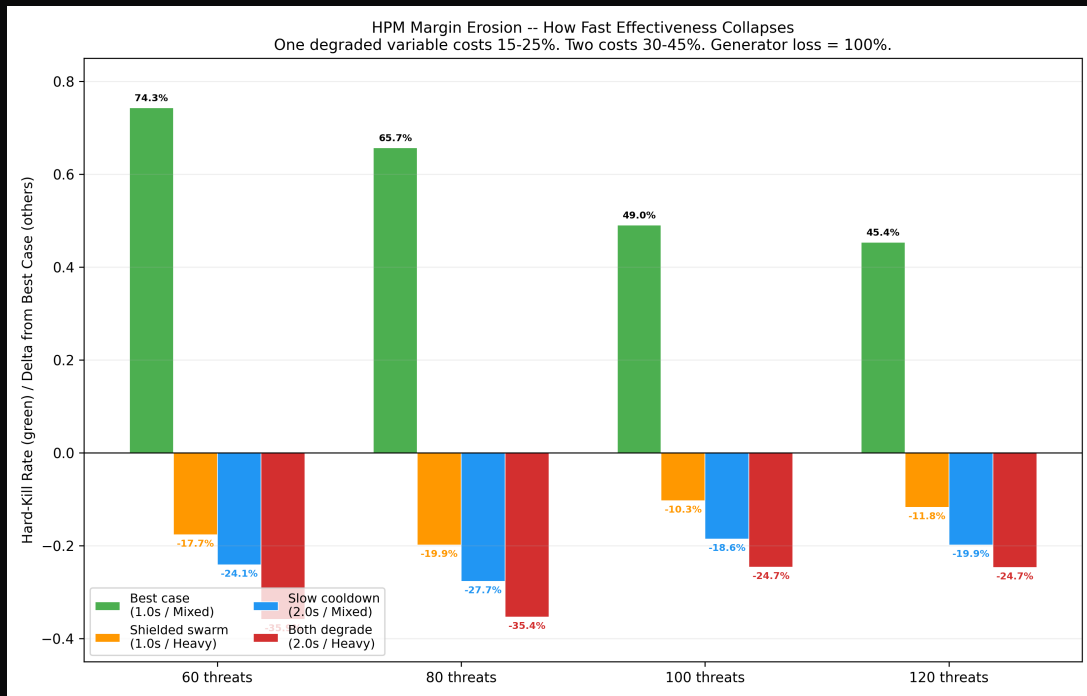


Figure 8: Margin erosion bars - how fast the kill rate collapses when variables degrade

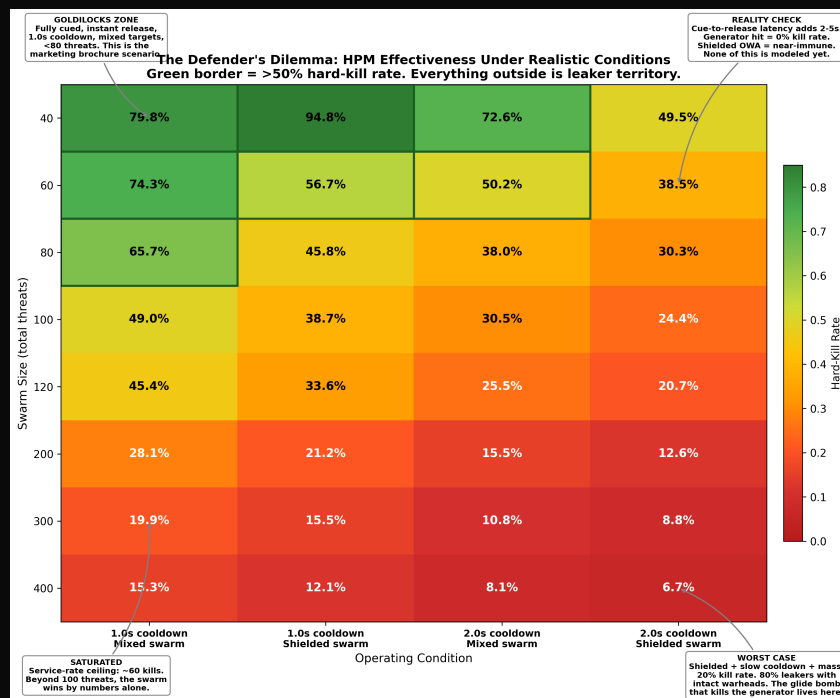


Figure 9: The Defender's Dilemma - full 8x4 envelope with operational reality callouts



## Results - Closed-Loop Cascade Failure

The upper-bound analysis answers "what can the HPM do?" The closed-loop analysis answers "what happens to the HPM when it does it?"

All 44 configurations were re-run with the simulation's damage engine routing warhead detonations from fixed-wing glide trajectories against defended assets. When a delta-wing OWA is hard-killed (flight controller latch-up), the airframe glides on its last heading. The armed warhead detonates at the glide terminal point. Blast overpressure and fragmentation are resolved against the nearest asset - the Leonidas, its sensors, or its generator.

### The Cascade Mechanism

1. HPM fires at Geran-2 at close range (saturated scenario, engaging late-arriving threats)
2. Flight controller latches up from cumulative gate-oxide stress (hard kill)
3. Delta-wing airframe glides on last heading at ~3:1 ratio with no thrust
4. Armed 50 kg warhead detonates on ground impact near defended asset
5. Blast and fragmentation damages or destroys sensors and, in some seeds, the Leonidas itself
6. If Leonidas destroyed: all subsequent HPM engagements stop - remaining inbound threats are unengaged leakers

The re-engagement mechanism - cumulative multi-pulse damage converting mission kills to hard kills - is itself the source of the cascade. The HPM's success at hard-killing targets creates the debris that damages the defensive system.

### Finding 7: Asset Survivability Under Cascade Failure

Density	Leonidas Kills (total / per seed)	Sensor Kills (total / per seed)	Asset Damage Events (per seed)
40	0 / 0.00	0 / 0.00	0.00
100	0 / 0.00	0 / 0.00	0.00
200	1 / 0.02	14 / 0.28	0.38
300	2 / 0.04	54 / 1.08	0.80



Density	Leonidas Kills (total / per seed)	Sensor Kills (total / per seed)	Asset Damage Events (per seed)
400	2 / 0.04	72 / 1.44	0.98

Table 4: Closed-loop cascade damage by swarm density (N=50 seeds, mixed-baseline composition)

**The good news: the Leonidas is remarkably survivable.** At 400 threats in the mixed-baseline configuration, the HPM was destroyed in only 2 of 50 seeds - a 96% survival rate. Even in the worst-case shielded configuration (gf\_400t\_1s\_sh), where engagement geometry forces more close-range kills against hardened targets, the Leonidas survived in 45 of 50 seeds - 90% survivability.

**The bad news: sensor attrition is the real cascade.** At 300+ threats, the engagement produces an average of 1+ sensor kills per seed. At 400 threats in the worst-case shielded configuration, sensors are destroyed at a rate of 1.76 per seed. Losing one or two of five sensors does not immediately blind the fire-control loop, but it degrades coverage geometry, introduces azimuth gaps, and reduces track quality - all of which compound in a multi-axis engagement.

**As noted, we only modeled one sortie.** Secondary sorties will exploit their advantage, whether they know they have it or not, with the blind spots in the sensor coverage. Of course, in the real world this is counteracted with overlapping sectors-of-coverage and dispersing your depth meaningfully.

## Hard-Kill Rate Delta

The closed-loop hard-kill rates are nearly identical to the upper-bound case across most configurations:

Density	Upper-Bound Hard-Kill Rate	Closed-Loop Hard-Kill Rate	Delta
40	79.8%	79.6%	-0.2%
60	74.3%	74.5%	+0.2%
80	65.7%	63.0%	-2.7%
100	54.5%	53.9%	-0.6%
200	28.1%	27.6%	-0.5%
300	19.9%	19.9%	0.0%
400	15.3%	15.2%	-0.1%

Table 5: Upper-bound vs closed-loop comparison (mixed-baseline density sweep)



The delta is  $\leq 3\%$  across all density levels. The HPM performs nearly identically whether or not asset damage is modeled - because in 96% of seeds, the Leonidas survives. In the 4% of seeds where the Leonidas is destroyed, the engagement ends and remaining threats become unengaged leakers. But across the full seed population, this contributes only a small reduction in the mean hard-kill rate.

The real impact of the cascade is not in the kill rate statistics. It is in the tail risk: the 4-10% of engagements where the HPM is destroyed mid-run and the defense collapses entirely.

**Known divergence - wave timing configs.** The wave timing configurations (Finding 4) show near-zero delta between upper-bound and closed-loop cases. This is expected: staggered waves allow the HPM to service each small group completely at range before the next arrives, which means hard kills occur at longer standoff distances where the glide terminal point falls well outside the defended perimeter. The cascade mechanism requires close-range kills in saturated scenarios - precisely the condition that staggered waves prevent. Both sweeps ran wave configs independently and the absolute kill counts converge within normal CI bounds (e.g., 77.7 vs 77.0 destroyed for the 5-wave config). The wave timing finding holds in both cases and is self-reinforcing: staggered waves help the defender *and* eliminate the cascade failure mechanism simultaneously.

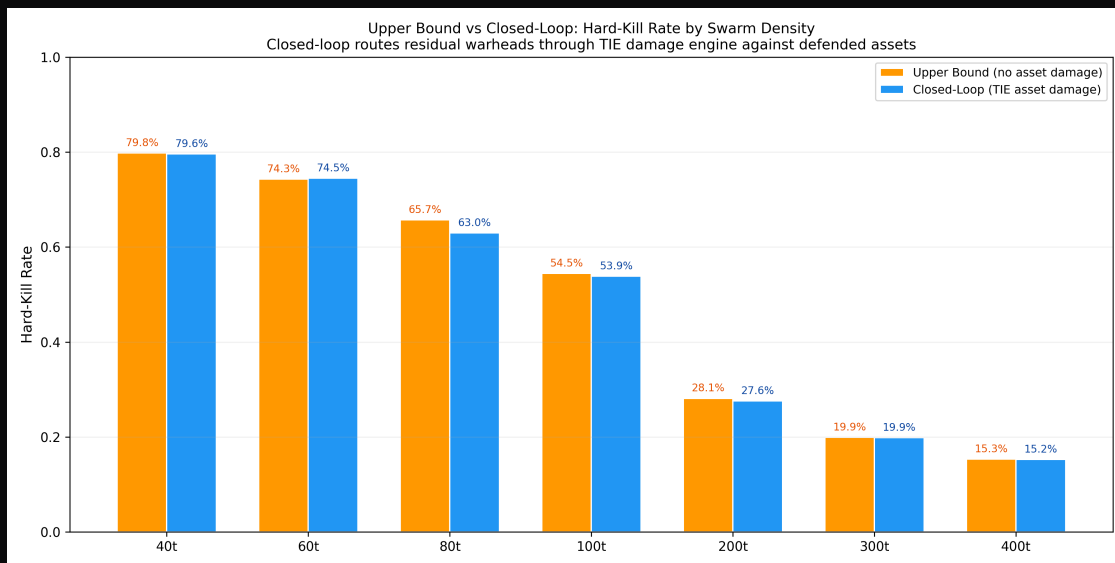


Figure CL-Delta: Upper-bound vs closed-loop hard-kill rate comparison by density

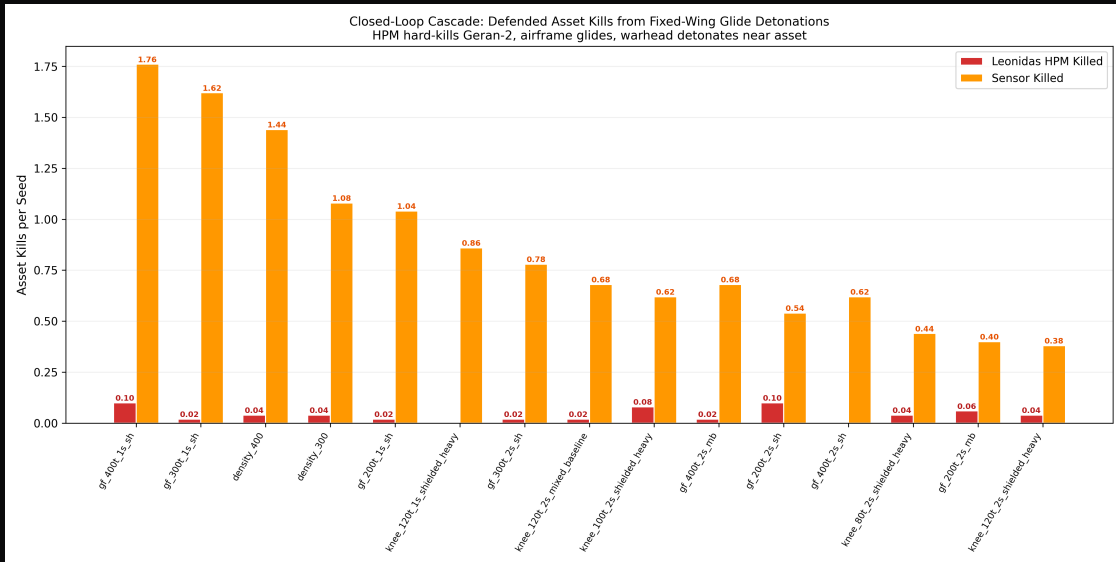


Figure CL-Asset: Leonidas and sensor kills per seed by configuration

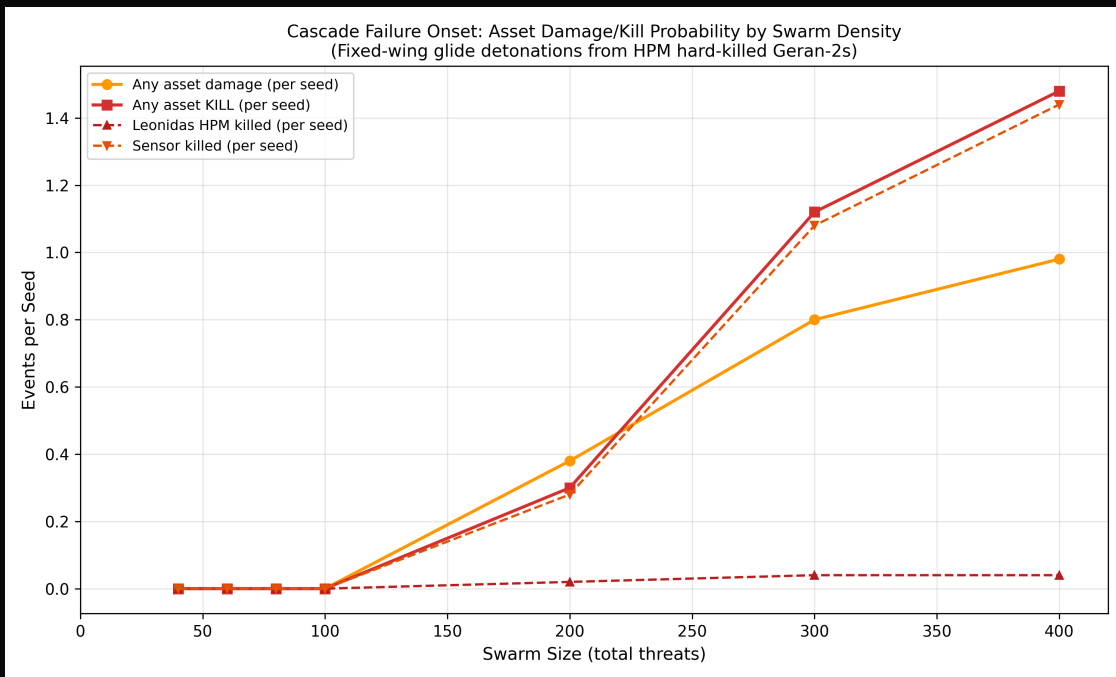


Figure CL-Probability: Cascade failure onset by swarm density

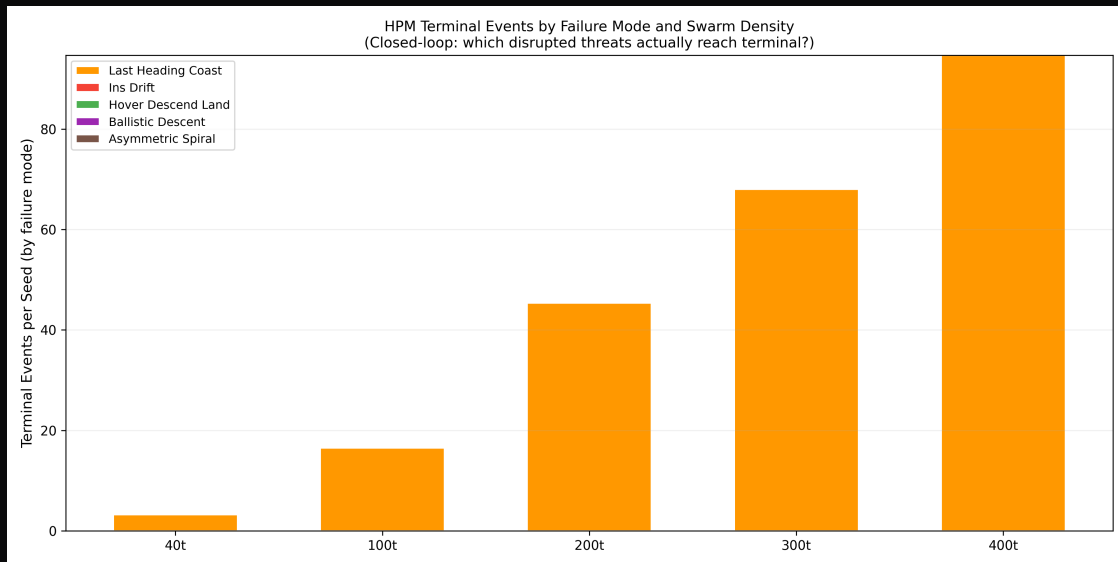


Figure CL-Terminal: Terminal events by failure mode

## Remaining Limitations of the Closed-Loop Model

The closed-loop model advances fidelity but is not complete. The following simplifications remain:

**Quadcopter debris is not modeled.** FPV and commercial multirotor hard kills produce a near-vertical crash (lawn-dart) below the kill point. Debris trajectory is not computed. This is acceptable - quadcopter debris is localized below the kill point with no significant horizontal displacement toward the defended asset.

**Single-target blast resolution.** Warhead blast and fragmentation are resolved against the single nearest asset, not as an area-of-effect across multiple nearby entities. In reality, a 50 kg HE detonation 30m from the Leonidas would also damage sensors at 50m - the model underestimates co-located damage.

**No secondary effects.** LiPo fires, smoke obscuration, sympathetic detonation of nearby UXO, and environmental degradation from blast are not modeled. These would further degrade the defense in saturated scenarios.

---

## Observations: HPM Value Beyond the Hard-Kill Rate

The headline findings regarding service-rate saturation, shielding dominance, residual warhead hazard, and cascade failure do not mean HPM is without value. The system has properties that the hard-kill rate alone does not capture.



## The Hype Is Real

There are plenty of articles from defense journalists singing the praises of HPM in dealing with swarms. While we did not doubt the coverage, it is more impressive seeing it simulated with the available physics. Real-world conditions can definitely influence our findings - it's hard to get access to detailed Foreign Instrumentation Signals Intelligence (FISINT) in the public domain - but what we have been able to model is nothing short of impressive. Having several dozen kills in a relatively short period of time, even with our modest time delay barely accounting for mechanically steering a single HPM, is very hard not to get excited over.

Against quadcopter strike drones - whether they're RF-, cellular-, or fiber-optically driven - it's absolutely devastating. Until such a time that collaborative, AI-driven drone swarms are proliferated along with the logistics support of setting up 60+ drones for a flight, it's hard to imagine dedicating what would be a company-sized element of expert drone aces to attack one area. That is firmly within the realm of OWAs, so it makes this even more impressive. Even against a flight of 6 drones, as long as you have quality Counter-UAS sensing, HPM will make short work of them and turn them into write-offs.

## Cost Exchange

At ~\$0.02 per pulse and ~60 pulses per engagement window, the HPM expends approximately \$1.20 in RF energy to service a full swarm engagement. Even accounting for generator fuel, maintenance, and amortized system cost, the cost-per-engagement is orders of magnitude below kinetic alternatives. A single Stinger missile costs approximately \$119,000. A single ESSM costs \$1.8 million. The HPM's per-shot economics are in a different universe.

Against a 100-threat swarm of \$500-\$20,000 drones, the HPM achieves 54.5 kills at a marginal cost of effectively zero per kill. The kinetic alternative - Phalanx CIWS at \$30 per round, 200 rounds per engagement - costs \$6,000 per swarm but has limited magazine depth and cannot service 54 simultaneous targets. The HPM's area-effect beam cone and near-zero marginal cost make it the only economically sustainable solution against mass drone warfare.

## Behavioral Effects as Force Multiplier

Even targets the HPM does not hard-kill are degraded. At 400 threats, the HPM produces 312.7 behavioral effects per seed - GPS disruptions, camera damage, navigation degradation. These effects convert precision strikes into area-effect hazards, reducing the probability that any individual drone achieves its intended aimpoint. Against a coordinated strike on a specific high-value target (radar array, command post, ammunition storage), converting the first wave from precision to area suppression can be the difference between a mission kill and a near miss.



## Illumination for HEL Inner Layer

The HPM's electronic effects produce observable behavioral signatures - trajectory changes, speed deviations, heading drift - that should make HPM-degraded targets easier for a downstream HEL to acquire and track. A drone executing a blind terminal dive has a predictable, uncontrolled trajectory. A drone drifting on INS has a smooth, linear flight path with no evasive jinking. These degraded targets are ideal HEL engagements: predictable geometry, known speed, no countermeasures.

## Discussion

### Electronic Defeat Is Not Hazard Elimination

The central finding of this research is not centered purely on kill rates, but on outcomes. The per-subsystem damage model reveals a spectrum of HPM effects that binary "kill/no-kill" models obscure:

Outcome	Definition	Warhead Status
Hard kill (multirotor)	FC/ESC/BMS destroyed, flight terminated	Lawn-dart crash below kill point, warhead status depends on fuze arming
Hard kill (fixed-wing)	FC/ESC/BMS destroyed, stabilization lost	Aerodynamic glide on last heading (~3:1 ratio), armed warhead detonates on ground impact
Mission kill	GPS/camera disrupted, cannot complete intended strike	Airframe intact, flies failure trajectory (INS drift, blind dive), armed warhead on degraded path
Guidance-defeated residual hazard	Guidance degraded, airframe intact, impacts within hazard radius	Armed munition on ballistic/drift trajectory

The gap between "mission kill" and "hard kill" is where the warhead lives. A mission-killed Geran-2 is not likely to hit its intended target, but it's still going to hit *something*. The HPM did its job - the precision strike was defeated - but the 50-kilogram warhead is still flying, still armed, and the airframe is still aerodynamically stable. Where it lands is now a probability distribution, not a certainty, and that distribution is centered on the defended asset.

The closed-loop analysis adds another dimension: a hard-killed fixed-wing OWA is not debris - it is a guided glide bomb aimed at whatever it was pointed at when the flight controller died. At close range in saturated scenarios, the glide terminal point is inside the defended perimeter. The HPM's own success creates the debris that threatens the HPM.



**Author's Note:** While we can model wind and other atmospheric conditions that could blow a fixed-wing drone on a terminal glide off course, that was not the intent of this research, and as noted we modeled perfectly calm and clear conditions.

Given that ordnance is tumbling or gliding out of the sky, this reframes the HPM (plus layered HEL, SHORAD, or other effectors) architecture from redundancy to consequence management:

Layer	Role
HPM outer perimeter	Disrupts and thins swarm. Area denial. Degrades guidance on OWA platforms.
HEL inner perimeter	Physically kills guidance-disrupted residuals. Destroys airframe and detonates/deflagrates warhead at altitude. Converts armed munitions into debris.
Kinetic point defense	All-weather backup and last-ditch hard kill. Operates when smoke/weather degrades HEL.
EOD / base response	Manages UXO and environmental aftermath. Clears impact zone.

The HEL's role is not just "kill HPM leakers." It is "destroy the warhead before it reaches the ground." A hard-killed multicopter falls as localized debris. A hard-killed fixed-wing glides as an armed munition. A mission-killed drone drifts as an armed munition. The HEL converts both of the latter into the former.

### The Shielding Arms Race

The attacker's natural evolution toward better shielding is the single biggest strategic threat to HPM dominance. Russia's drone forces went from commercial quadcopters to carbon-fiber FPVs to Geran-series composite OWA in approximately 18 months. Each iteration increased effective shielding. The next step - metal-tube fuselages, EMI gaskets, deliberately hardened wiring - is a predictable engineering response to HPM deployment.

At 8 dB shielding (Geran-5 class surrogate, based on OSINT and open-source FISINT extrapolation), the HPM's per-pulse kill probability drops below meaningful levels at tactical ranges. The cumulative multi-pulse model can still accumulate damage, but only if the fire-control loop can allocate multiple engagement cycles per target. It is unlikely a single HPM can do so when the swarm is large enough to saturate the service rate, or is part of a simultaneous, multi-axis attack.

The defender has the following options, in no particular order:

- Higher peak power, which requires proportionally larger apertures or closer ranges
- Focused beam mode to provide a narrower beamwidth and higher gain at the cost of a smaller engagement cone



Multi-effector coordination, and/or multiple HPMs covering different sectors

Layered HEL/kinetic backup for shielded targets

All of these increase cost, complexity, and logistics footprint - exactly the asymmetry that favors the attacker in drone warfare economics.

## What This Analysis Does Not Address

This study tests one HPM effector defending one fixed asset against one swarm. The findings should not be interpreted as "HPM C-UAS is insufficient." They should be interpreted as "single-effector HPM C-UAS has a narrow operating envelope whose boundaries are quantifiable." The broader defensive architecture includes capabilities not modeled:

**Multi-effector HPM coverage.** Two or four Leonidas-class emplacements covering different sectors would extend the engagement cone to 120 deg-360 deg and multiply the service rate proportionally. This changes the saturation arithmetic.

**HEL inner layer.** A 30-150 kW HEL engaging HPM-degraded residuals at 1-2 km range would convert mission kills into hard kills. The behavioral effects from HPM exposure (predictable trajectory, no evasive maneuver) make HPM-degraded targets ideal HEL engagements. Likewise, in favorable weather conditions, the HEL can engage at slightly longer standoff ranges depending on power output and targeting sensors available.

**Kinetic point defense.** Phalanx CIWS, SHORAD, or man-portable air defense for most-weather backup when smoke, dust, and/or weather degrades the HEL. This "golden triangle" should be the localized, minimum viable defense architecture as part of a larger IADS/BMD effort.

**EW and deception.** GPS spoofing, command-link jamming, RF decoys, remote takeover, cyber soft-kill, and Military Deception (MILDEC) to complement HPM effects.

**Left-of-launch.** Dedicating Intelligence, Surveillance, Reconnaissance, and Targeting (ISRT) assets to speed up the F3 (Find, Fix, Finish) loop against UAV control points, UAV logistics hubs, and fixed-site OWA batteries (e.g., such as those in Donetsk Airport).

---

## Forward Vector

The data points toward four capabilities that the physics demands:

**Speed-of-light early kill & cleanup.** The HEL converts (or bypasses) HPM mission kills into hard kills by physically destroying the airframe and warhead at altitude. The cost per shot is measured in single-digit dollars.



The JLWS program (150-300 kW containerized, FY2027 Navy R&D) is directly responsive to this requirement, and the Army's IFPC-HEL (300 kW class) addresses the higher end. The engineering challenge is atmospheric propagation and dwell time against hardened targets - solvable problems on a known trajectory.

**Multi-effector coordination.** A single HPM defending a single point is the wrong unit of analysis. Two to four coordinated emplacements with sector coverage and handoff protocols multiply the service rate and close the geometric coverage gap. The sensor fusion and fire-control integration is the hard problem - not the emitter physics.

**Adaptive shielding intelligence.** The fire-control loop needs real-time shielding classification to prioritize engagement order. Killing FPVs first (correct behavior for maximizing total kills) is wrong behavior when the shielded Geran-5 with 50 kg HE is the existential threat. Classification from EO/IR PID, RF signature, radar cross-section, and behavioral profile should drive engagement priority, not detection order.

**Consequence management architecture.** If the HPM is deployed near populated areas, the residual warhead hazard - UXO, fire, contamination - requires pre-planned response. EOD teams, fire suppression, civilian evacuation routes, and medical response must be co-located with the defensive system. The HPM does not just defend an asset; it creates a debris field that must be managed. The closed-loop cascade findings make this even more urgent: glide detonations from hard-killed fixed-wing OWAs produce blast damage inside the defended perimeter even when the HPM is performing optimally.

---

## Methodology

### Simulation Engine

The Empyrean Defense Decision Dominance Engine Wargaming & Simulation Cyber Range module, running in batch mode via internal research CLI. The engine resolves per-subsystem HPM damage, behavioral failure trajectories, terminal impact adjudication, and - in the closed-loop configuration - warhead blast and fragmentation damage against defended assets using Kingery-Bulmash overpressure polynomials and Gurney fragment velocity models. All runs are deterministically reproducible from scenario file and seed.

### Two-Case Methodology

The full 44-configuration sweep was executed twice:

**Upper-bound case.** Hard-killed entities do not produce glide trajectories or warhead detonations. Defended assets are invulnerable. This establishes the theoretical ceiling of HPM performance - the number reported in



vendor demonstrations and product literature.

**Closed-loop case.** Hard-killed fixed-wing OWAs produce aerodynamic glide trajectories based on their airframe type, speed, altitude, and heading at the moment of kill. Armed warheads detonate at the glide terminal point. The damage engine resolves blast overpressure and fragmentation against the nearest defended asset. If the Leonidas or a sensor is destroyed mid-engagement, subsequent engagements are affected accordingly. This establishes the operationally realistic performance accounting for self-inflicted debris.

Both cases use identical stochastic seeds for direct comparison. The delta between cases isolates the cascade failure contribution.

### Stochastic Uncertainty Model

Five uncertainty dimensions applied per run under deterministic seed control:

Dimension	Distribution	Parameters
Polarization	LogNormal	sigma = 0.15, clamp [0.1, 1.0]
Coupling attenuation	Gaussian in dB	Card-specific spread (0.5-2.0 dB)
E50/sigma multipliers	LogNormal	sigma = 0.10, clamp [0.5, 2.0]
Spawn timing	Gaussian	+/-2s (1 sigma)
Spawn position	Gaussian	+/-25m (1 sigma)

When the stochastic mode is disabled, all perturbations are zeroed and the model reverts to deterministic behavior (Phase 1 compatibility). Same seed always reproduces exactly.

### Convergence Verification

The baseline density\_100 configuration was run at N=50 and N=250 seeds:

N	Hard-Kill Rate	CI Half-Width
50	0.545	0.024
250	0.641	0.013

Six configurations with CI half-width exceeding 0.03 at N=50 received Phase 2 refinement at N=250. All refined configurations converged below the 0.03 target. CI bounds are computed via percentile bootstrap (B=5,000 replicates, deterministically seeded for reproducibility).



## Statistical Note: Deterministic vs Stochastic

The Phase 1 parametric sweeps were fully deterministic - all seeds produced identical results because the HPM engagement pipeline had zero stochastic components. Phase 2 introduced five stochastic uncertainty dimensions. The stochastic results reported in this paper use Phase 2 methodology throughout. Phase 1 deterministic results are available for comparison but are not reported as primary findings.

## Author's Note: The Warhead That Keeps Flying

There is a version of this paper that would have been easier to write. The one where we report the Leonidas hard-kills 80% of a 40-drone swarm and declare HPM the future of counter-UAS. The demo reel version. The one that gets me on stage at AUSA (maybe, probably not anymore!).

That version is true! Alas, it is also incomplete.

The complete version includes the 8 drones the HPM disrupted but didn't destroy. It includes the Geran-2 that lost GPS at 1,200 meters, switched to INS, drifted 2 degrees over 22 seconds, and put a 50-kilogram warhead 82 meters from the thing it was defending. It includes the FPV that lost its camera feed, executed a blind terminal dive at 30 degrees plus or minus 10, and hit the ground 173 meters from the operator's last known aimpoint - which was the radar.

And then there's the version we didn't expect to write - the one where the HPM hard-kills a Geran-2 at close range, the delta-wing glides 120 meters on its last heading, and the warhead detonates 3 meters from the Z-Scan sensor. In 5 of 50 seeds at maximum density, that detonation took out the Leonidas itself. The HPM's own success created the debris that killed the HPM.

The disrupted drone doesn't know it's been disrupted. It doesn't know the GPS is gone. The INS says everything is fine. The autopilot maintains heading, maintains altitude, maintains speed. The airframe is aerodynamically stable. The warhead is armed. The only thing that changed is where it lands - and "where it lands" went from deterministic to stochastic. The probability distribution is centered on the defended asset because that's where the drone was pointed when you pulled the guidance.

**This is the sentence I keep coming back to: the HPM displaces the point of impact from the intended target to a stochastic dispersion zone. It does not eliminate the warhead.** Trust me, I love the idea (and videos) of quadcopters falling out of the sky, but they fall *somewhere*.



In the 40-threat scenario - the one where HPM is dominant, where the kill rate is 80%, where the vendor demo looks great - an average of 4.8 warhead impacts per engagement land within 500 meters of the Leonidas. At 80 threats, it's 7.1. These are not edge cases. These are the mean across 50 stochastic seeds with coupling uncertainty, polarization perturbation, and spawn jitter. The close-in hazard is a feature of the physics, not a bug in the model.

The layered defense argument is not about redundancy. It is about what happens to the warhead after the guidance dies. The HPM kills the mission. The HEL kills the airframe. The kinetic point defense kills what the HEL misses. And the EOD team handles what gets through everything - because something always gets through everything. A colleague once described directed energy as "the physics of making the other guy's day worse." That is accurate. The HPM makes the drone's day dramatically worse. But the drone doesn't have a day. It has a warhead, and the warhead doesn't care about its day.

That warhead? Perhaps it has purpose-built secondary fuzing designed to kill the EOD technicians and demining teams who respond to the debris. Perhaps a focused beam destroys the GPS of an OWA or a MALE drone at 4,000m AGL and it glides into someone's apartment building, their garden, a school, a hospital?

In civilian self defense courses, of the many things drilled into your head, the one that stands out the most is that you're responsible for where the bullet goes, physics be damned. There is a sort of inversion here, and it is not the fault of a soft-kill ground based air defense (SKGBAD) crew, but that warhead flies off somewhere. I'd love to say that it matters not, it'll just self-destruct or the point-detonating fuze will cause it to explode. What if it doesn't?

This is a research paper, but if you haven't gotten the drift, I value human life - which is why I take the time to collect my thoughts at the end.

As an American, I want our guys and gals to come home safe, I want HPM arrays damn near everywhere to take out UAS and turn them into lawn darts. That said, I also don't want a young Ukrainian boy, or young Russian girl, or an elderly man or woman to stumble upon a 300g grenade from a bomber drone nor a 50kg warhead from an OWA and be vaporized. That is to say nothing of the secondary and tertiary hazards. Your HPM array can soft-kill a flight of drones and they still set your LNG terminal ablaze. They set the oil fields on fire and send acid rain and crude oil downstream, and downwind.

That is why residual-risk accounting belongs in the kill chain, not in the footnotes.

It's easy not to care because physics don't care, they just *are*. Warheads don't care, HPMs don't care, they're not human. We are. Let's not forget it.



- Empyrean 7

---

## References

### Primary Physics Sources

1. Friis, H.T., "A Note on a Simple Transmission Formula," *Proc. IRE*, 34(5), pp. 254-256, 1946.
2. Balanis, C.A., *Antenna Theory: Analysis and Design*, 4th ed., Wiley, 2016.
3. Backstrom, M. and Lovstrand, K.G., "Susceptibility of Electronic Systems to High-Power Microwaves: Summary of Test Experience," *IEEE Trans. EMC*, 46(3), pp. 396-403, 2004.
4. Jafari, H. and Anbarjafari, G., "Directed Energy Weapons Against Unmanned Aerial Vehicles: A Comprehensive Review and Modeling Framework," *Defence Technology*, 2026.
5. Zhang, Y. et al., "GPS Interference Effects on Commercial UAS Navigation Systems Under High-Power Microwave Exposure," *J. Electromagnetic Waves & Applications*, 39(4), 2025.

### Standards

6. MIL-STD-461G, "Requirements for the Control of Electromagnetic Interference Characteristics of Subsystems and Equipment," US DoD, 2015.
7. IEC 61000-2-13, "Electromagnetic compatibility (EMC) - Part 2-13: Environment - High-power electromagnetic (HPEM) environments - Radiated and conducted," 2005.
8. IEEE Std 145-2013, "IEEE Standard for Definitions of Terms for Antennas," 2013.
9. ITU-R P.676-13, "Attenuation by atmospheric gases and related effects," ITU, 2022.
10. ITU-R P.838-3, "Specific attenuation model for rain for use in prediction methods," ITU, 2005.

### Patent Sources (Leonidas Surrogate Parameters)

11. US Patent 11,616,295 B2, "Adaptive high-power microwave system," Epirus Inc., 2023.
12. US Patent 11,658,410 B2, "Modular high-power microwave system," Epirus Inc., 2023.
13. US Patent 12,068,618 B2, "High-power microwave beam steering," Epirus Inc., 2024.

### Statistical Methods

14. Efron, B. and Tibshirani, R.J., *An Introduction to the Bootstrap*, Chapman & Hall, 1993.
15. Marsaglia, G., "Xorshift RNGs," *J. Statistical Software*, 8(14), 2003.
16. Aumasson, J.-P. et al., "BLAKE2: simpler, smaller, fast as MD5," ACNS 2013. RFC 7693.



---

## OSINT and Operational Sources

17. Epirus press release, "Leonidas defeats 49-drone swarm in live demonstration," September 2025.
  18. DJI, "Failsafe Behavior Documentation," DJI Developer Portal, 2024.
  19. Analog Devices, "ADIS16480 Tactical Grade IMU Datasheet," Rev. C, 2023.
-



## Appendix A: HPM Physics Model - Complete Formula Catalog

All formulas below are implemented in the Wargaming & Simulation Cyber Range module.

### A.1 Free-Space Power Density

$$S(R) = (P_{\text{peak}} \times G) / (4\pi \times R^2) \quad [W/m^2]$$

Symbol	Definition	Units
P_peak	Peak pulse power	W
G	Antenna gain (linear, from dBi)	dimensionless
R	Slant range to target	m

Source: Friis transmission equation (far-field). H. T. Friis, "A Note on a Simple Transmission Formula," Proc. IRE, 1946.

Validation: Leonidas surrogate at 1 km:  $S = (450 \times 10^6 \times 31.62) / (4\pi \times 1000^2) = 1,131 \text{ W/m}^2$ .

### A.2 Antenna Gain Conversion

$$G_{\text{linear}} = 10^{(G_{\text{dBi}} / 10)}$$

Source: Standard dB-to-linear conversion. IEEE Std 145-2013.

### A.3 Atmospheric Attenuation

$$\begin{aligned} \tau_{\text{atmo}}(R) &= \tau_{\text{clear}}(R) \times \tau_{\text{rain}}(R) \times \tau_{\text{humidity}}(R) \\ S_{\text{effective}} &= S(R) \times \tau_{\text{atmo}}(R) \end{aligned}$$

Clear-air: ITU-R P.676-13 frequency-dependent absorption. Rain: ITU-R P.838-3  $k \times R_{\text{rate}}^{\text{alpha}}$ . Water vapor: ITU-R P.676-13 resonance lines.

At L-band (1.1 GHz), atmospheric attenuation is negligible (< 0.01 dB/km).  $\tau \sim 1.0$  for ranges < 5 km.

### A.4 Coupling Attenuation

$$S_{\text{coupled}} = S_{\text{incident}} \times 10^{(-SE_{\text{dB}} / 10)}$$

Source: Shielding effectiveness per MIL-STD-461G, adapted for incidental drone airframe shielding.

Airframe	SE_dB	Basis
Foam/plastic (DJI Mini)	0	No conductive material



Airframe	SE_dB	Basis
Carbon fiber frame (FPV racer)	4	CF is conductive, partial Faraday cage
Fiber-optic FPV (CF frame)	5	CF + no RF antenna apertures
Composite (Shahed-136 class)	2	Fiberglass/foam, unshielded wiring
Metal tube fuselage (Geran-5 class)	8	Al profiles + steel structure (OSINT)
Military UAS with EMI gaskets	15-30	MIL-STD-461G compliant enclosures

### A.5 Polarization Mismatch

$$\eta_{pol} = \cos^2(\phi)$$

For random target aspect:  $\langle \cos^2\phi \rangle = 0.5$  (3 dB average loss). Floor at 0.1 for cross-polarization coupling paths. Per-entity deterministic seeding ensures reproducibility; stochastic perturbation applied in Phase 2.

Source: Jafari & Anbarjafari (2026), section 4.6.

### A.6 Peak Electric Field

$$|E| = \sqrt{S \times \eta_0} \quad [V/m]$$

Where  $\eta_0 = 377 \text{ ohm}$  (free-space impedance). Valid in far-field ( $R > 2D^2/\lambda$ ).

Validation: At 1 km,  $S = 1,131 \text{ W/m}^2 \rightarrow |E| = 653 \text{ V/m}$ . At 3 km,  $S = 126 \text{ W/m}^2 \rightarrow |E| = 218 \text{ V/m}$ .

### A.7 Far-Field Validity Gate

$$R_{ff} = 2D^2 / \lambda$$

Leonidas surrogate:  $D = 2.4 \text{ m}$ ,  $f = 1.1 \text{ GHz}$ ,  $\lambda = 0.273 \text{ m}$ .  $R_{ff} = 42.2 \text{ m}$ . All tactical engagements ( $> 100 \text{ m}$ ) are well into the far field.

### A.8 Per-Subsystem Sigmoid Damage

$$P_{kill}(|E|) = 1 / (1 + \exp(-(|E| - E_{50}) / \sigma_E))$$

Subsystem	E50 (V/m)	sigma_E (V/m)	S50 (W/m <sup>2</sup> )	Failure Mode
GPS/GNSS LNA	150	30	~60	Front-end burnout -> nav loss
Camera/FPV	200	40	~106	Image sensor saturation/damage



Subsystem	E50 (V/m)	sigma_E (V/m)	S50 (W/m <sup>2</sup> )	Failure Mode
Flight Controller	250	50	~166	CMOS latchup -> loss of stabilization
ESC (motor control)	300	60	~239	Gate oxide breakdown -> motor stop
BMS (battery)	350	70	~325	Protection circuit trip -> power loss

Source: Calibrated to Backstrom & Lovstrand (2004) CMOS latchup thresholds and Zhang et al. (2025) GPS LNA susceptibility.

Validation at 3 km ( $|E| = 218$  V/m): GPS  $P = 0.91$ , Camera  $P = 0.62$ , FC  $P = 0.34$ , ESC  $P = 0.08$ , BMS  $P = 0.02$ . This produces the desired graduated kill - GPS and camera likely damaged, FC marginal, ESC and BMS survive. Multiple pulses needed for hard kill.

## A.9 Cumulative Multi-Pulse Hazard

```
P_survive(N) = prod(i=1..N) (1 - P_kill_i)
P_damage(N) = 1 - P_survive(N)
```

Each pulse is an independent stochastic stress trial. NOT thermal accumulation - voltage-driven damage for nanosecond pulsed HPM. Source: Jafari & Anbarjafari (2026), section 4.14.

Validation: At 2 km,  $P_{\text{per\_pulse}} = 0.15$  for ESC. After 5 pulses: 0.56. After 10 pulses: 0.80.

## A.10 Correlated Subsystem Evaluation

Shared coupling factor computed ONCE per target per pulse (incident field x coupling attenuation x polarization mismatch), then each subsystem's sigmoid evaluated against that shared coupled stress. If coupling is favorable, ALL subsystems are more likely to fail. If poor, ALL survive.

## A.11 Sidelobe Fratricide

```
G_sidelobe = G_main - SLL_dB [dBi]
S_sidelobe(R) = (P_peak x G_sidelobe_linear) / (4pi x R^2)
P_fratricide = sigmoid(S_sidelobe / S_threshold)
```

At 200m from a representative radar LNA (30 W/m<sup>2</sup> susceptibility threshold):  $S_{\text{sl}} = 895$  W/m<sup>2</sup> -> certain damage. At 2,000m:  $S_{\text{sl}} = 8.95$  W/m<sup>2</sup> -> safe. Minimum safe standoff approximately 1,500-2,000m for radar-class sensors.

## A.12 Failure Trajectory Models



**Blind terminal dive (camera loss on FPV):** Dive angle drawn from Gaussian distribution ( $\mu=30$  deg,  $\sigma=10$  deg), heading jitter from Gaussian ( $\mu=0$  deg,  $\sigma=15$  deg). At 100m AGL, 30 m/s: mean ground distance 173m, 95% impact ellipse  $\sim 80\text{m} \times 90\text{m}$ .

**INS drift (GPS loss on autonomous fixed-wing):** Heading drift 1-3 deg/min (commercial MEMS IMU). Speed and altitude maintained. The airframe continues flying on its pre-disruption velocity vector with slowly accumulating heading error. Source: ADIS16480 datasheet.

**Hover-descend-land (commercial GPS failsafe):** Decelerate, hover 3s, descend at 1.5 m/s, land. Source: DJI failsafe documentation.

**Fixed-wing glide (FC/ESC hard kill on delta-wing OWA):** When the flight controller latches up on a fixed-wing airframe (Geran-2, Geran-5), the delta-wing has inherent aerodynamic stability. It does not tumble - it glides on its last heading at approximately 3:1 glide ratio with no thrust. The armed warhead detonates on ground impact. Ground distance from kill point = speed  $\times$  (altitude / descent\_rate). At 1,000m AGL and 46 m/s, glide distance is approximately 2,000m. At 60m AGL (close-range kill in saturated scenario), glide distance is approximately 120m - placing the warhead detonation within the defended perimeter. This is the mechanism that produces the closed-loop cascade failure.

**Multicopter ballistic descent (ESC kill on quadcopter):** Terminal velocity  $\sim 20$  m/s (2-3 kg,  $C_d \sim 1.0$ ). Near-vertical crash below kill point. No significant horizontal displacement. Warhead fuze arming state depends on impact velocity and fuze type.

**Asymmetric spiral (partial ESC failure):** Heading drift 30-90 deg, descent 8-12 m/s, spiral radius 50-100m.

---

## Appendix B: Stochastic Uncertainty Model

All stochastic perturbations are deterministically seeded. Same seed reproduces exactly. Different seeds produce plausible alternative realizations of uncertain physical parameters.

### B.1 PRNG

xorshift64 (Marsaglia 2003) for uniform variates, Box-Muller transform for Gaussian. Identical algorithms across both compiled and interpreted components for cross-language reproducibility.

### B.2 Polarization Perturbation

```
multiplier = exp(sigma_pol * z) [LogNormal, sigma=0.15]
```



$\eta_{pol\_perturbed} = \eta_{pol\_deterministic} \times \text{clamp}(\text{multiplier}, 0.1, 1.0)$

68% of realizations within [0.86x, 1.16x] of nominal. Source: Jafari & Anbarjafari (2026), section 4.6.

### B.3 Coupling Attenuation Perturbation

$SE_{perturbed} = SE_{nominal} + \text{gauss}(0, \text{spread\_dB})$   
 $SE_{perturbed} = \text{clamp}(SE_{perturbed}, 0, 2 \times SE_{nominal})$

Target Class	SE_nominal	Spread	Rationale
Commercial FPV	4 dB	1.0 dB	Variable CF layout
Fiber-optic FPV	5 dB	1.0 dB	Similar variance
Geran-2	2 dB	0.5 dB	Consistent construction
Geran-5	8 dB	2.0 dB	Variable metal/composite, OSINT uncertainty

### B.4 E50/Sigma Multipliers

$\text{mult\_i} = \exp(\text{sigma\_e50} \times z\_i)$  [LogNormal, sigma=0.10, clamp [0.5, 2.0]]  
 $E50_{perturbed} = E50_{nominal} \times \text{mult\_i}$

68% within [0.90x, 1.11x] of nominal. Clamp bounds consistent with Backstrom & Lovstrand (2004) observed 1.5-3x variance across CMOS device lots.

### B.5 Spawn Jitter

Timing: +/-2s (1 sigma). Position: +/-25m (1 sigma). Defender assets excluded from jitter. Wave ordering preserved.

### B.6 Bootstrap CI

Percentile method, B=5,000 replicates, alpha=0.05. Convergence target: CI half-width < 0.03. Deterministically seeded for reproducibility.

## Appendix C: Average Power Budget

$P_{avg\_rf} = P_{peak} \times \tau \times PRF = 450 \text{ MW} \times 20 \text{ ns} \times 50 \text{ Hz} = 450 \text{ W}$   
 $P_{wall} = P_{avg\_rf} / \eta = 450 / 0.30 = 1,500 \text{ W}$



$$P_{\text{waste}} = P_{\text{wall}} \times (1 - \eta) = 1,050 \text{ W}$$

At 1.0s engagement cycle, cooling rate (~3.0 kJ/s) outpaces heat generation (~2.2 kJ/pulse). The thermal model does not limit engagement duration in tactical-length engagements. The real vulnerability is geometric (60 deg beam cone), tactical (multi-axis saturation), and logistical (generator survivability) - not thermal.

---

*This analysis uses physics-grounded simulation models with OSINT-derived surrogate parameters. It is not an operational assessment. Results represent system-class behavioral characterization under stated assumptions with stochastic uncertainty bounds. The authors have no affiliation with any defense contractor or government agency referenced herein.*

UNCLASSIFIED // OPEN SOURCE INTELLIGENCE ONLY



## EMPYREAN DEFENSE

Decision Dominance Through Sensor Fusion

<https://empyreandefense.com/>

This document is UNCLASSIFIED and contains no controlled or restricted information. All data sources are open-source and publicly available. Distribution is unlimited.

UNCLASSIFIED // OPEN SOURCE INTELLIGENCE ONLY