

Dark Space Detection

Inline Enforcement Against Traffic Directed at Non-Existent Network Addresses

Francesco Trama

Founder and CEO, PacketViper LLC, Pittsburgh, Pennsylvania, USA

ORCID: 0009-0004-8437-6351

Technical Capability Brief

Version 1.0, September 2026

License: CC BY 4.0 DOI: 10.5281/zenodo.22284418

Abstract

Dark Space Detection is the automatic identification of, and inline enforcement against, any traffic directed at a network address where no legitimate host exists. If the set of live addresses on a segment is known at any given moment, any packet destined for an address outside that set is a reliable indicator of reconnaissance or lateral movement. No decoy has to be touched and no signature has to match; the destination itself is the tell. This brief describes how PacketViper implements the capability using an existing sensor paired with an inverted context group of active hosts, how it relates to automated moving target defense (AMTD) and deception responders, why it closes the residual coverage gap identified in the PacketViper AI agent containment test, and why operational technology (OT) segments, with small and stable address populations, are the strongest fit for the approach.

Keywords: dark space detection, network reconnaissance, lateral movement, automated moving target defense, deception, inline enforcement, operational technology, autonomous AI agents, preemptive cybersecurity

1. What It Is

Dark Space Detection is the automatic identification and enforcement against any traffic directed at a network address where nothing legitimate exists.

If the set of IP addresses alive on a segment is known at any given moment, any packet destined for an address outside that list is an immediate red flag. No decoy has to be touched. No signature has to match. The traffic itself is the tell.

PacketViper implements this natively using an existing capability: a sensor paired with an inverted context group listing all active IPs on the segment. Everything outside that list is dark space. Any probe into dark space triggers automatic inline enforcement.

2. How It Works

1. PacketViper maintains a live map of active hosts on the segment through ARP and traffic observation.
2. A context group is built from that live map. This is the known, legitimate IP population.
3. The sensor is configured to invert that context group. This arms every address that is not a real host.
4. Any traffic destined for a non-existent address fires enforcement instantly, inline, with no human in the loop.

The attacker does not have to be unlucky enough to hit a deceptive responder. They only have to go somewhere nothing lives. In any real network, that happens on the first probe.

3. Where It Plays

3.1 Within AMTD

AMTD continuously rotates the deceptive surface: IPs, ports, banners, services. Dark Space Detection extends that principle to the gaps between the decoys. An attacker cannot map around what is deceptive if even the empty space is armed.

The attack surface is not just moving. It is fully covered. Static plus dynamic plus dark space leaves no safe ground to probe.

3.2 Within Deception Responders

Traditional deception requires the attacker to touch a deceptive responder to trigger detection. That is a probabilistic bet. The more decoys, the better the odds, but the probability never reaches 100 percent.

Dark Space Detection removes that dependency. Detection no longer requires contact with a specific false service. Contact with any non-existent address is sufficient. The coverage ceiling moves from "depends on decoy density" to "anything that probes where nothing lives."

3.3 Within Sensors

Sensors are the PacketViper inline enforcement engine. A sensor with an inverted context group is the implementation of Dark Space Detection: no new architecture, no added complexity. It is a sensor configuration that transforms known-good IP space into a boundary condition. Everything outside that boundary is automatically treated as hostile.

Sensors already detect and act. Dark Space Detection gives them a new trigger: not only "touched a decoy" but also "went somewhere that should not exist."

4. Why It Matters for AI Agents

The PacketViper AI agent containment test identified a theoretical 3 percent gap: a sufficiently patient agent that somehow mapped and avoided every deceptive responder might reduce the probability of deception contact.

Dark Space Detection closes that gap architecturally.

An autonomous AI agent performing reconnaissance must probe IP addresses to discover hosts. In any network, the vast majority of the address space is unoccupied. The agent does not need to touch a deceptive responder to be detected. It only needs to probe an address where nothing exists, which is statistically certain during any reconnaissance sweep.

There is no reconnaissance strategy that avoids both deceptive responders and dark space simultaneously, because dark space is everywhere decoys are not.

5. Why It Matters for OT

OT environments are ideal for Dark Space Detection. The IP population is small, stable, and well known. A 20-device OT segment has roughly 20 live addresses in a /24. The remaining 234 addresses are permanently dark.

Any attacker, whether human, malware, or autonomous AI agent, performing lateral movement inside that segment will probe addresses that do not exist before finding the ones that do. Every one of those probes is a detection event.

No endpoint agents. No protocol-specific logic. No prior knowledge of the attacker's tools. The network itself becomes the detector.

6. The One-Line Version

If you touch something that does not exist, we already know you are there.

Related Work by the Author

This brief extends the following Zenodo records in the same series:

- [1] Trama, F. (2026). Denying the World Model: Automated Moving Target Defense as an Architectural Countermeasure to Autonomous AI Agents. Zenodo. <https://doi.org/10.5281/zenodo.21347479>
- [2] Trama, F. (2026). A Packet-Layer Capability Triad for Remote Operational Technology: A Candidate Compensating Control for Physically Exposed, Resource-Constrained Sites. Zenodo. <https://doi.org/10.5281/zenodo.21403201>
- [3] Trama, F. (2026). Detection Is Not Containment: A 168-Hour Model of Alert-to-Enforcement Latency in IT and OT. Zenodo. <https://doi.org/10.5281/zenodo.22100402>

Conflict of Interest Disclosure

The author is the founder and CEO of PacketViper LLC, the vendor of the product described in this brief. The capability is described by its observable effect and configuration so that the underlying claim, that traffic destined for a non-existent address is a reliable and enforceable indicator, can be evaluated and tested independently of any specific product.

License

This work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>.