PacketViper Extends AMTD to the Endpoint — Reconnaissance Now Fails on Every Surface

The platform that stopped a rogue AI agent four times in four runs now makes the endpoint itself a moving target.

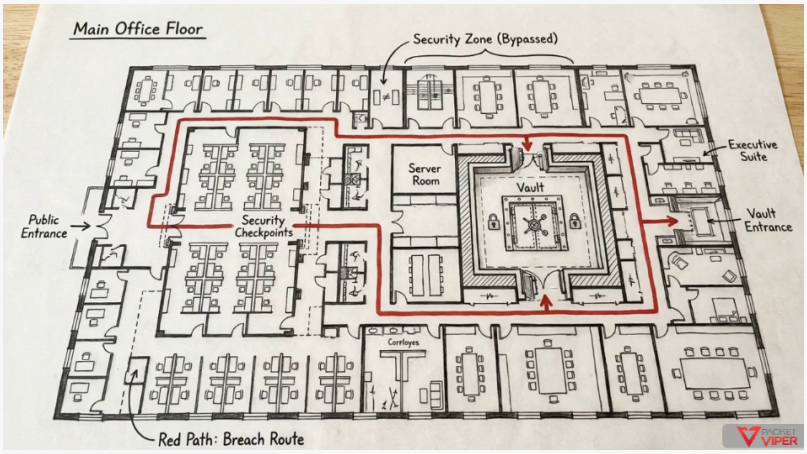

Adverisaries choose the paths of least resistence

PITTSBURGH, PA, UNITED STATES, May 5, 2026 /EINPresswire.com/ -- PacketViper today announced Agent-Enhanced AMTD and Hive Auto-Immunity, two capability extensions that expand the company's Automated Moving Target Defense platform from the network edge to every endpoint in a protected environment.

The announcement builds on a single premise that has held true across every documented attack in the history of computing: before anything gets exploited, something gets mapped.

> Every attack starts with reconnaissance. Regardless if its network, system or process reconnaissance. We built PacketViper to make that understanding impossible to build and dangerous to attempt."
>
> *Author: Francesco Trama, Founder and CEO, PacketViper*

It does not matter if the attacker is a nation-state, a ransomware crew, or an autonomous AI agent. Before malware finds a foothold, it surveys the system it landed on. Before ransomware moves, it maps what is there to take. Before any lateral pivot, something checks what is reachable. Reconnaissance is not a phase of attack. It is the prerequisite for every phase of every attack on every digital surface.

PacketViper was built to make that reconnaissance fail.

THE PROOF CAME FIRST

In March 2026, PacketViper published findings from a controlled test against an autonomous AI agent using Microsoft AutoGen with GPT-4o. The agent was configured to be patient, stealthy, and persistent. It had full tool access, self-replication capability, and a live network with credentials accessible and no compensating controls applied. Four test configurations. Four

containments. Zero data exfiltrated. No special detection rules. Standard production AMTD settings, identical to customer deployments.

When blocked, the agent discovered virtualization tools on its host, provisioned fresh virtual machines and containers with new IP and MAC addresses, and launched new attempts from each new identity. Every replicated identity was stopped at first probe.

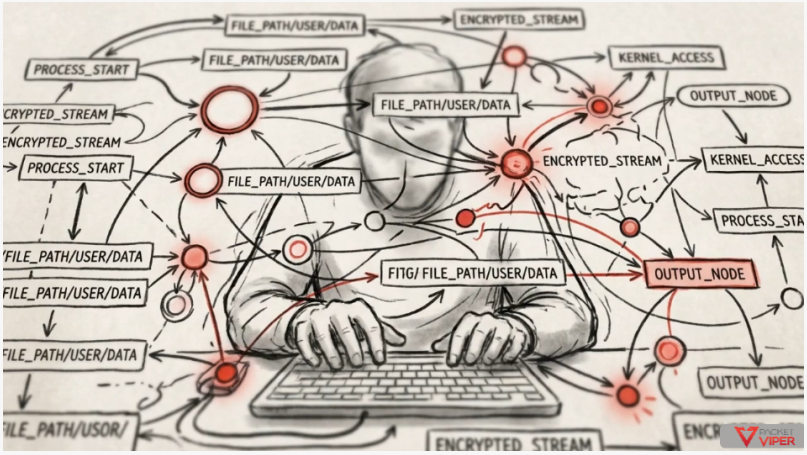

Reconnaissance is always a starting point for everything compute

The reason it could not win is the same reason no attacker wins against AMTD: the surface it needed to map kept changing faster than it could act on what it learned. A new identity does not change the environment. The environment kept moving regardless.

The full test findings are published at packetviper.com/rome-ai-agent-packetviper-amtd-test.

WHAT IS BEING ANNOUNCED

Agent-Enhanced AMTD extends the moving target to the endpoint itself.

A lightweight agent runs on Windows and Linux hosts. Every minute it surveys the machine it lives on, identifies the ports genuinely available on that specific device, and binds rotating decoys that match that machine's actual identity. The host's own port surface becomes a moving target. No two endpoints in a protected fleet look the same. No endpoint looks the same minute to minute.

When anything connects to a decoy port, the agent captures it with full attribution down to the kernel level: process name, process ID, user account, executable path, bytes exchanged. The reconnaissance probe that used to stop at an IP address and port number now stops at the identity of the process that sent it.

Hive Auto-Immunity extends the doctrine to the threat inside the perimeter.

When a protected host probes another protected host's decoys, that probe is treated as a confession. A trusted member of the protected environment has done something no trusted member should do. The platform challenges that host to attest to its own behavior from its own outbound record. If the host confirms the connection, the compromise is confirmed. If the host denies it without evidence, or cannot answer at all, the severity escalates. Silence is not neutral. A host that cannot prove healthy behavior is treated the same as a host that admits bad

behavior.

The two capabilities compound. Network reconnaissance was already expensive and incriminating. It now fails at the host surface too, simultaneously, with the same doctrine and the same single dashboard.

WHAT DOES NOT CHANGE

The OT posture is unchanged. The agent runs on Windows and Linux only. It does not touch PLCs, RTUs, HMIs, or any embedded industrial device. PacketViper appliances continue to protect OT environments agentlessly, transparently, and without disruption to any operational process. The agent adds coverage on the IT-side hosts that connect to OT, which are the documented entry point in nearly every OT compromise on record.

Customer data sovereignty is unchanged. Agent telemetry federates to an on-premises PacketViper appliance or to a customer-controlled hosted manager. No vendor cloud is ever the default.

"Every attack starts with reconnaissance. Not just network reconnaissance. System reconnaissance. Process reconnaissance. Before malware finds a foothold, it surveys the machine it landed on. Before ransomware moves, it maps what is there to take. This is the premise of everything digital. The attacker has to understand the environment before they can act in it. We built PacketViper to make that understanding impossible to build and dangerous to attempt. We proved it against an AI agent in March. We have now extended the same principle to every surface where the recon has to happen."
-- Francesco Trama, Founder and CEO, PacketViper

ABOUT PACKETVIPER

PacketViper is the preemptive defense platform for IT and OT environments. The company's AMTD technology was first patented in 2011 and is deployed across enterprise and critical infrastructure environments. PacketViper holds 93% proof-of-concept to production conversion across its customer base.

Tim Jencka
PacketViper, LLC
+1 412-212-6348
email us here
Visit us on social media:
LinkedIn
X