Stream.Security

Autonomous Production Defense

One Live Model Of Production. No Boundaries.

Stream is the AI-native platform built to fight AI-enabled attacks. It autonomously prevents, detects, hunts, and remediates exposures and threats across production at machine speed - driving risk toward zero. It replaces the fragmented stack of scanners, runtime agents, exposure tools, and playbooks with one live model of production.

Gartner Cool Vendor
For Modern InfoSec

4.8 / **5** Customer Rating

9.6 / **10** CSAT

Attackers Don't Think In Production Boundaries Anymore

Cloud, SaaS, identity, runtime, AI, network and on-prem no longer behave as separate estates. They operate as one connected system, and attackers move through them as one. The security industry answered with the opposite structure: separate tools, separate consoles, separate collection cycles - scanning for misconfigurations, listing vulnerabilities, producing findings for people to validate by hand. Adding AI agents on top of that structure only produces findings faster. **The deeper problem is what a finding is.** A finding describes the environment as it was when the scan ran. An attack happens in the environment as it is now. Finding a misconfiguration or a vulnerability is no longer the work. Closing exploitability, safely, inside a system that is still running, is the work.

Defending Production Needs A New Approach

Stream is the only Autonomous Production Defense Platform that works across your entire production estate. It runs on a patented **ProductionTwin®**, a high-fidelity security data harmonization layer that models cloud, SaaS, identity, runtime, AI, network, perimeter, on-prem VMware, security controls, and the behavior running on top of them into one live model of production: **real-time, fully correlated, continuously updating. Not a snapshot**. And it does not stop at boundaries - the boundaries that fragment every other tool are the same boundaries an attacker moves across. Inside ProductionTwin they are one system.

Why Is Autonomy Possible Here? And Why Can You Trust It?

Autonomy is not a model problem. It is a ground-truth problem - an agent is only as safe as what it reads before it acts. Rebuild context out of logs and last night's scan, and the action is a guess about the present dressed as a decision - so a human is inserted and becomes the bottleneck that made autonomy pointless. Read a live model and every step is computed instead: reachability, blast radius, impact before execution, closure after. The model enables the AI. The AI does not enable the model.

Four Autonomous Loops. One Live Model.

Each runs on its own, continuously, with no one starting it.

Autonomous Prevention

Closes exploitable paths on its own, before anyone must defend them.

Stream finds the paths through production that are exploitable and closes them itself: tightening access and permissions, enforcing boundaries and guardrails, correcting configuration, reverting drift. It simulates every change against the live model first, so it can act without waiting for a sign-off.

Autonomous Detection

Follows the attack across every boundary, with no analyst driving it.

Stream weights every detection by live reachability and blast radius, not by a static score. On its own it follows one identity across the entire production estate and ephemeral compute as one chain, tunes its own detection and plants canaries on the paths an attacker is likeliest to take next.

Autonomous Threat Hunting

Hunts live production continuously, with no one starting the hunt.

Identity and attack-path hunts run against the environment as it is now, correlated across boundaries and reconstructing workloads that lived for minutes. Whatever a hunt proves, Stream promotes into a permanent detection or canary by itself.

Autonomous Remediation and Response

Contains the attack and remediates the vulnerability, without causing the outage.

Stream computes blast radius live, generates the containment options and the remediation flows, and simulates each against production before anything executes. It then acts at the least-impact point - perimeter, identity, network, control or workload - and recomputes the path to prove it is gone. All of it unprompted; you decide where approval is required.

How Would Your Current Production Defense Deal With An AI-Powered Attack?

A modern attack, through the Hugging Face lens 🤗

The attack ran across four boundaries. One identity moved through all four.

Correlating that means manual pivoting across four telemetry sources and four retention windows: no tool holds the chain together, and no analyst pivots fast enough to catch one that is already gone. And that was the slow version: an AI-driven chain crosses the same boundaries before a responder finishes authenticating into the second tool.

The Attacker's Clock — **Minute 0 to 156**

01 Endpoint

A classic foothold. Your EDR fires and the ticket gets closed.

The Gap
The story continues where EDR does not look.

02 Cloud

Short-lived tokens carry it into cloud workers.

The Gap
The token expired before the next scan ran.

03 SaaS and Git

Reaches the code repo. Different console, different owner, different login.

The Gap
Nobody joins a repo event to a loud event.

04 Ephemeral Compute

Spins up Lambda and K8s workloads, uses them, tears them down. Lived three minutes.

The Gap
It never entered an inventory. You cannot query a resource that is gone.

Attack Complete — Minute 156

Your clock starts at minute 157 - not because your team is slow, but because no product owns the chain. It has to be assembled by hand across four consoles, and that work cannot begin until someone notices, long after the attacker has finished.

On Stream, there is no second clock.

One chain from minute 0, blast radius computed as it moves, containment simulated before it executes, closure proven after.

The Road to Preemptive Production Defense

Yes, you have a CNAPP, SaaS security, ITDR, a firewall, a SIEM, and an AI SOC - but then what?

A critical CVE is published. Thousands of workloads report it, and the severity score says the same thing about all of them. That is where a CNAPP stops. Preemptive defense starts one question later. Stream answers six questions instead - autonomously, against production as it is at that second.

01 Can It Be Reached?	02 What Sits Behind It?	03 How Can The Path Be Broken?
Whether an attacker can actually get to this workload, and from where.	Which identities, data and systems fall with it if it does.	Every option that removes exploitability, not only the patch.
04 What Would Each Option Cost?	**05 Which Point Is Safest?**	**06 Is It Actually Closed?**
Each one simulated against live production before anything runs.	Perimeter, identity, network, security control or the workload itself.	The path is recomputed. If it still computes, it is not closed.

Questions 01–03 need a live model, and questions 04–06 are impossible without one - because you cannot simulate an action against a snapshot, and you cannot prove a path is closed by re-reading last night's scan. Stream answers all six itself, computing reachability and blast radius, generating and simulating the options, executing the fix at the least-impact point and verifying closure against the live model, so the risk is shut without a ticket, without an analyst, and without a maintenance window to find out whether it worked.

Comprehensive and Extensible

Every Stream capability is a mission running against ProductionTwin. Prevention, detection, hunting and response are not four engines - they are four standing missions, and the platform is programmable, so your team writes its own. **StreamForce** is Stream's agent harness: you define the goal, the instructions, the tools an agent may use, and the data it may reach.

It inherits the live model, so there is no pipeline to build and no context to feed it. Agents compose with each other through shared context, simulate every action against ProductionTwin before executing, and leave the data, the reasoning, and the approvals traceable. Building a new defense stops being a quarter of engineering and becomes a mission definition.

The hard part was never the automation. It was having something true enough to automate against.

AAA · PENNYMAC · OneMain Financial · cross river · TC Energy · naf new american FUNDING · deluxe · RingCentral

> "Most weeks now, the first I hear about an incident is that Stream already closed it. My team isn't hunting anymore - we're reviewing what it already found and shut down."

Mike Young
Director of Cybersecurity, Risk & Compliance

Hunt Energy

> "I used to triage alerts all day. Now I read a list of things that are already closed. That's the entire difference."

Arye Shulman Ehrenreich
CIO at Shield

shield.

> "Time used to be the currency of cloud. Now we don't spend it on attacks that never happen - Stream closes the exploitable path before anyone ever gets to use it."

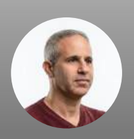

Tamir Ronen
CISO at HiBob

HiBob