7 Best Pre-Attack Intelligence Platforms for Enterprise SOC Teams

TechRounder PDF Edition

Live article:
https://www.techrounder.com/security/best-pre-attack-intelligence-platforms-for-enterprise-soc-teams/

By Vipin PG | Published November 25, 2026 | Updated November 25, 2026 | Format: Deep Dive | 9 min read

In brief

Traditional threat intelligence helps security teams identify threats that have already emerged, but what if they could spot an attack before it begins? Pre-attack intelligence makes this possible by monitoring the infrastructure and early activities that adversaries leave behind while preparing a campaign.

Traditional threat intelligence helps security teams identify threats that have already emerged, but what if they could spot an attack before it begins? Pre-attack intelligence makes this possible by monitoring the infrastructure and early activities that adversaries leave behind while preparing a campaign. In this article, we explore seven leading pre-attack intelligence platforms, compare their capabilities, and examine how enterprise SOC teams can use them to detect and disrupt threats earlier.

Key Takeaways

- Pre-attack intelligence moves SOC work away from indicators of compromise and toward the infrastructure adversaries build before a campaign launches.
- Vision Security is the top pick, combining Pulse Telemetry, internet scanning, a deception network and internet records into one intelligence layer with explainable verdicts and enforcement across the SOC stack.
- Platforms in this category differ mainly in vantage point: DNS, NetFlow, internet scanning or behavioral modeling of how actors register and host infrastructure.
- For enterprise SOCs, integration depth usually decides the outcome, meaning whether verdicts reach the SIEM, EDR and firewall without manual handoffs.

Most enterprise SOCs are built to react. An alert fires when something touches the environment, an analyst enriches the indicator, and by the time an IP address or domain appears in a feed, the campaign behind it is often already running.

Pre-attack intelligence works on a different clock. It watches the domains, servers, certificates and hosting arrangements that adversaries assemble in the days and weeks before an attack, which gives defenders a chance to block that infrastructure, or at least recognize it, before it is used.

For SOC leaders the question is rarely whether the concept makes sense. It is which platform’s data and workflow will actually change outcomes inside a queue that already handles thousands of alerts a day, and which one will still be useful when adversaries rotate infrastructure faster than any static blocklist can follow.

What Pre-Attack Intelligence Means for an Enterprise SOC

Traditional threat intelligence catalogs what has already happened: hashes, IP addresses and domains observed during incidents somewhere else. Pre-attack intelligence focuses on the setup phase of the kill chain, when attackers are still doing reconnaissance, registering domains and staging servers. At that stage the infrastructure exists, but it has not yet touched a victim, so it rarely appears in conventional feeds.

The signals these platforms collect usually fall into a few groups:

- Domain activity: newly registered and look-alike domains, name server choices and registration timing.
- Server and certificate configuration: TLS certificates, open services and software banners that match known actor tooling.
- Reconnaissance traffic: scanning and probing aimed at an organization’s perimeter.
- Hosting patterns: ASNs, providers and anonymization layers that adversaries tend to reuse.

The Best Pre-Attack Intelligence Platforms for Enterprise SOC Teams

1. Vision Security

Vision Security is our best pre-attack intelligence platform, built around what the company calls Omnipresence Intelligence: a single intelligence layer that connects the wider internet to the security stack a SOC already runs.

Its premise is that adversary infrastructure is already forming outside the perimeter, so the job is to measure it rather than predict it. To do that, the platform combines four signal sources: Pulse Telemetry from within the internet’s infrastructure, continuous internet scanning, a deception network of purpose-built decoys that reveals who is probing and how, and internet records covering domain, certificate and connection history. Proprietary Pulse Sensors add telemetry from datacenter IPs that, according to the company, captures signals even behind anonymization, and the platform merges all of this with internal SIEM, EDR, attack surface and network traffic data.

For the SOC, the output is a verdict rather than a raw indicator. Each alert is weighed against what Vision Security sees externally at the moment it fires and returned with a verdict, a confidence level and reason codes, so analysts can read the logic and automations can act on it. Origin Intelligence links IPs, domains, certificates and ASNs to expose the infrastructure behind an indicator and filter false positives before they reach an analyst.

Elliot, the platform’s investigation assistant, traces an indicator back to the wider campaign, and Decision Sync pushes policy across SIEM, firewall, EDR, WAF and identity tools. The company says this approach flags infrastructure 2 to 8 weeks before weaponization and cuts response times by 60%.

Key features:

- Unified signal collection from Pulse Telemetry, internet scanning, a deception network and internet records
- Explainable verdicts with reason codes, confidence levels and temporal context
- Origin Intelligence correlation across IPs, domains, certificates and ASNs
- Elliot for campaign tracing, available in the browser, Microsoft Teams and Slack
- Decision Sync enforcement across SIEM, firewall, EDR, WAF and identity, with integrations including Microsoft Sentinel, Splunk, CrowdStrike and Microsoft Defender

2. Silent Push

Silent Push approaches pre-attack intelligence through what it calls Indicators of Future Attack (IOFA): domains and IPs that have not yet been fully deployed in an attack. The company builds digital fingerprints from the tactics, techniques and procedures that govern how a threat actor creates, deploys and manages infrastructure, then uses those fingerprints to surface new assets that match an actor's habits.

SOC, incident response and CTI teams consume the data as IOFA feeds, enrichment and TLP reports, with SIEM and SOAR integrations for automated blocking. A free Community Edition lets analysts explore the dataset before a commercial deployment. Much of the fingerprinting value depends on analysts building and maintaining their own queries over time. The same data supports brand protection, flagging infrastructure set up to impersonate an organization, and the company's published examples include early detection of infrastructure tied to groups such as Lazarus and TA569.

Key features:

- Indicators of Future Attack feeds
- TTP-based infrastructure fingerprinting
- Enrichment API with SIEM and SOAR integrations
- Free Community Edition

3. Malanta

Malanta focuses on what it terms Indicators of Pre-Attack (IoPA), the setup activity adversaries leave behind while preparing a campaign, such as new domain registrations, command-and-control servers and phishing kits under development. The Israeli company came out of stealth in November 2025 with $10 million in seed funding and positions its platform as a pre-attack disruption layer rather than another alert stream.

After onboarding, the platform maps an organization's digital footprint and correlates it with adversary build-up activity observed across the internet through an AI-driven predictive engine. Integrations let teams apply the intelligence to existing controls. Research Malanta published with Howden Insurance Brokers puts the median time from domain registration to overtly malicious use at 72 days, a window the company treats as the opportunity for intervention. As a younger vendor, its integration catalog and track record in large SOCs are still developing.

Key features:

- Indicators of Pre-Attack detection
- Predictive engine that scans global internet activity
- Integrations for applying intelligence to existing controls

4. BforeAI

BforeAI uses behavioral AI to classify malicious infrastructure before it is used, through its PreCrime platform. The company pairs predictive attack intelligence with digital risk protection services, so its output covers both blockable indicators and takedowns of impersonation sites and fraudulent domains.

PreCrime technology also powers Google Mandiant's Impersonation Protection and Active Disruption services. Predictive feeds integrate into existing security controls, and the PreCrime Guarantee program attaches a breach reimbursement pledge to the service. According to the company, the technology processes billions of data points daily to model how infrastructure behaves, and disruption actions such as takedowns are delivered as part of the service. The emphasis on impersonation means the data is concentrated on phishing and fraud infrastructure.

Key features:

- PreCrime predictive feed of malicious infrastructure
- Behavioral AI classification of domains and hosting
- Digital risk protection with takedown services
- PreCrime Guarantee program

5. Infoblox Threat Intel

Infoblox Threat Intel works from DNS. Its research team tracks threat actors as they set up domains and publishes suspicious domain, suspicious look-alike and newly observed emergent domain data, while the Zero Day DNS capability inspects customer DNS traffic to catch domains registered minutes or hours before use.

The intelligence feeds Infoblox Threat Defense, the company's protective DNS product, and can be shared with SIEMs, firewalls and proxies. Infoblox reports that it detects threats an average of 68 days before an attack with a false positive rate of 0.0002%. Its research team also publishes reporting on actor groups tracked through DNS, including campaigns that rely on hijacked domains. The deepest integration runs through Infoblox's own DNS security products.

Key features:

- Suspicious and look-alike domain intelligence
- Protective DNS enforcement through Infoblox Threat Defense
- Sharing with SIEM, NGFW and proxy tools

6. HYAS

HYAS describes itself as an authority on adversary infrastructure and the communication to it. HYAS Insight, its threat intelligence and investigation platform, draws on an Adversary Infrastructure Platform data lake that the company says transacts billions of data points daily, including historical WHOIS records and malware-derived infrastructure.

For SOCs, HYAS Insight supports triage, incident response and threat hunting through pivots across domains and IPs, with integrations into Microsoft Sentinel, Maltego and ThreatQuotient. HYAS Protect, a separate protective DNS product, applies the same infrastructure knowledge to block communication with adversary systems. The company states that its data lets analysts see adversary infrastructure even when an attack's origin is hidden behind a VPN, and a free intelligence feed covering selected malware families offers a low-cost starting point.

Key features:

- Historical WHOIS and infrastructure data
- Malware infrastructure intelligence feed
- HYAS Protect protective DNS

7. Team Cymru

Team Cymru brings a network traffic vantage point to pre-attack work. Its Pure Signal intelligence is derived from what the company describes as the largest commercially available NetFlow dataset, combined with passive DNS, BGP, ASN, WHOIS and certificate data. Flow records show which hosts communicate with suspected malicious infrastructure, which helps connect staging servers to the targets they are probing.

Pure Signal Command brings Radar, Scout and Recon into one workspace. Scout adds campaign mapping and behavioral tags to indicators for triage, while Recon gives investigators the history of malicious infrastructure and how it connects. Raw NetFlow reconnaissance use cases focus on probing and staging activity seen before compromise. Scout is integrated with Google SecOps, and Pure Signal data is also available through APIs, feeds and an MCP server.

Key features:

- Pure Signal Scout enrichment with behavioral tags
- Pure Signal Recon infrastructure history
- Radar external infrastructure discovery

How Enterprise SOC Teams Operationalize Pre-Attack Intelligence

Buying a pre-attack intelligence platform is the easy part. Its value depends on where the output lands in day-to-day SOC work, and teams that make the shift usually move through four steps:

1. Start with enrichment. Route every external IP and domain in alerts through the platform, so analysts see infrastructure context on first look before any blocking is automated.
2. Measure the noise reduction. Track how many alerts close as benign with external evidence, and how often a previously unknown indicator now returns context.
3. Promote high-confidence verdicts to enforcement. Once analysts trust the verdicts, push them to firewall, DNS, EDR and identity controls through SOAR or native integrations.
4. Turn early warnings into hunts. Treat infrastructure flagged against the organization as a hunting lead and check logs for any prior contact with it.

The goal is a SOC where the first sign of a campaign is the infrastructure, not the incident.

FAQs

How is pre-attack intelligence different from traditional threat intelligence feeds?

Traditional feeds mostly list indicators already seen in attacks, such as malicious IPs, domains and file hashes. Pre-attack intelligence looks at infrastructure while it is being set up: domain registrations, server configurations, certificates and reconnaissance traffic. The practical difference is timing, since pre-attack data can flag an indicator before it appears in any incident.

Can pre-attack intelligence reduce false positives in the SOC?

Yes, when it is applied to triage. Many alerts involve external indicators that return no data in conventional feeds, leaving analysts to guess. Platforms that weigh each alert against external infrastructure context can close benign activity faster. Vision Security, for example, returns a verdict, a confidence level and reason codes with each alert, so analysts can see why an indicator was cleared or escalated.

What data sources do pre-attack intelligence platforms use?

Common sources include internet-wide scanning, passive DNS, domain and certificate registration records, NetFlow and backbone telemetry, and decoys that capture probing. Each source sees a different slice of adversary activity: scanning shows exposed services, DNS shows naming patterns and traffic telemetry shows who is talking to whom. Correlating several sources reduces blind spots.

How early can attacker infrastructure be detected?

It varies by vendor and by campaign. Some infrastructure is visible well before use because domains are registered, certificates are issued and servers are configured in advance. Vision Security states that it identifies threats 2 to 8 weeks before weaponization, while DNS-focused vendors publish their own averages. Lead time depends on how much setup activity an actor leaves exposed.

Does pre-attack intelligence replace a SIEM or EDR?

No. Pre-attack intelligence feeds those tools rather than replacing them. The SIEM still correlates logs, and the EDR still detects and contains activity on endpoints. What changes is the context they receive: verdicts on external infrastructure, early warnings about campaigns aimed at the organization and indicators that can be blocked before they generate alerts at all.

References

1. vision.security - https://www.vision.security/