ANOMALI

Solutions Brief

THREAT INTEL MODERNIZATION BRIEF

Modern SOCs take in an average of 4,484 alerts a day. Analysts spend close to three hours triaging them, and about 83% turn out to be false positives . Adding more indicators to that pipeline does not relieve the pressure.

For more than a decade, a threat intelligence program was judged by accumulation: feeds subscribed to, indicators ingested, reports published on a schedule. The number told you how much your team could take in. It said nothing about whether any of it changed a decision. That model assumed an analyst downstream would read the intel and work out the next move, which was a fair assumption when threats developed over weeks.

Indicators of compromise age fast. Infrastructure rotates, and adversaries build evasion of signature-based detection into their operations deliberately. A bulk IOC list is often stale by the time it reaches the person meant to act on it. Automated scanning has climbed to as high as 36,000 scans per second globally . Intelligence that answers “have I seen this IP before” is answering a question the attacker has already moved past.

WHERE THE VALUE ACTUALLY SITS

The more useful test is what a piece of intelligence lets you decide. Three questions matter to the person on shift: is this activity real or noise, is it relevant to my environment, and what do we do next. When intelligence answers those in the moment, it stops being a reference document and becomes decision support. Intelligence that can’t act is just expensive reporting.

Getting there depends on context the indicator alone cannot carry. Behavior and campaign context, how activity maps to known tactics and whether it belongs to a larger operation, holds up even when a specific indicator is brand new or short-lived. That is what lets a team catch something earlier and rank it correctly instead of waiting on a match against a list. There is a structural side to it as well. Intelligence produced in a system that sits apart from the SOC has to be carried across by hand, and that handoff is exactly where relevance leaks out. The direction of travel is intelligence that reaches detection, prioritization, and investigation on its own: who, why, and what next, not just what happened.

QUESTIONS WORTH ASKING OF YOUR CURRENT PROGRAM

If you are trying to tell whether your intelligence supports decisions or mostly describes threats, a few things separate the two:

- How long does it take for a piece of intelligence to change something real, a control, a detection, a priority? If the answer is measured in days, that tells you something important.
- Does intelligence reach detection and investigation by itself, or does someone move it there manually every time?
- When an indicator is novel, does the program have anything to say, or does it go quiet until a signature catches up?
- Can you reconstruct, after the fact, why a given action was taken? Decision support you cannot explain is hard to stand behind.

None of this requires ripping out what you run today. The move from indicator lists to decision support is a change in what you expect of intelligence before it is a change in tooling. Teams that make the shift early tend to catch threats sooner and spend less of the shift on noise.

Anomali Managed Intelligence as a Service, powered by ThreatStream Next-Gen, fuses attacker and campaign context into detection, prioritization, and investigation, so a brand-new indicator still gives your team something to act on instead of going quiet until a signature catches up. See how it works.

SEE ANOMALI IN ACTION

Request a Demo