Trellix

Five Dimensions of Operational Threat Intelligence for Machine-speed Defense

Building the threat intelligence advantage in the age of AI

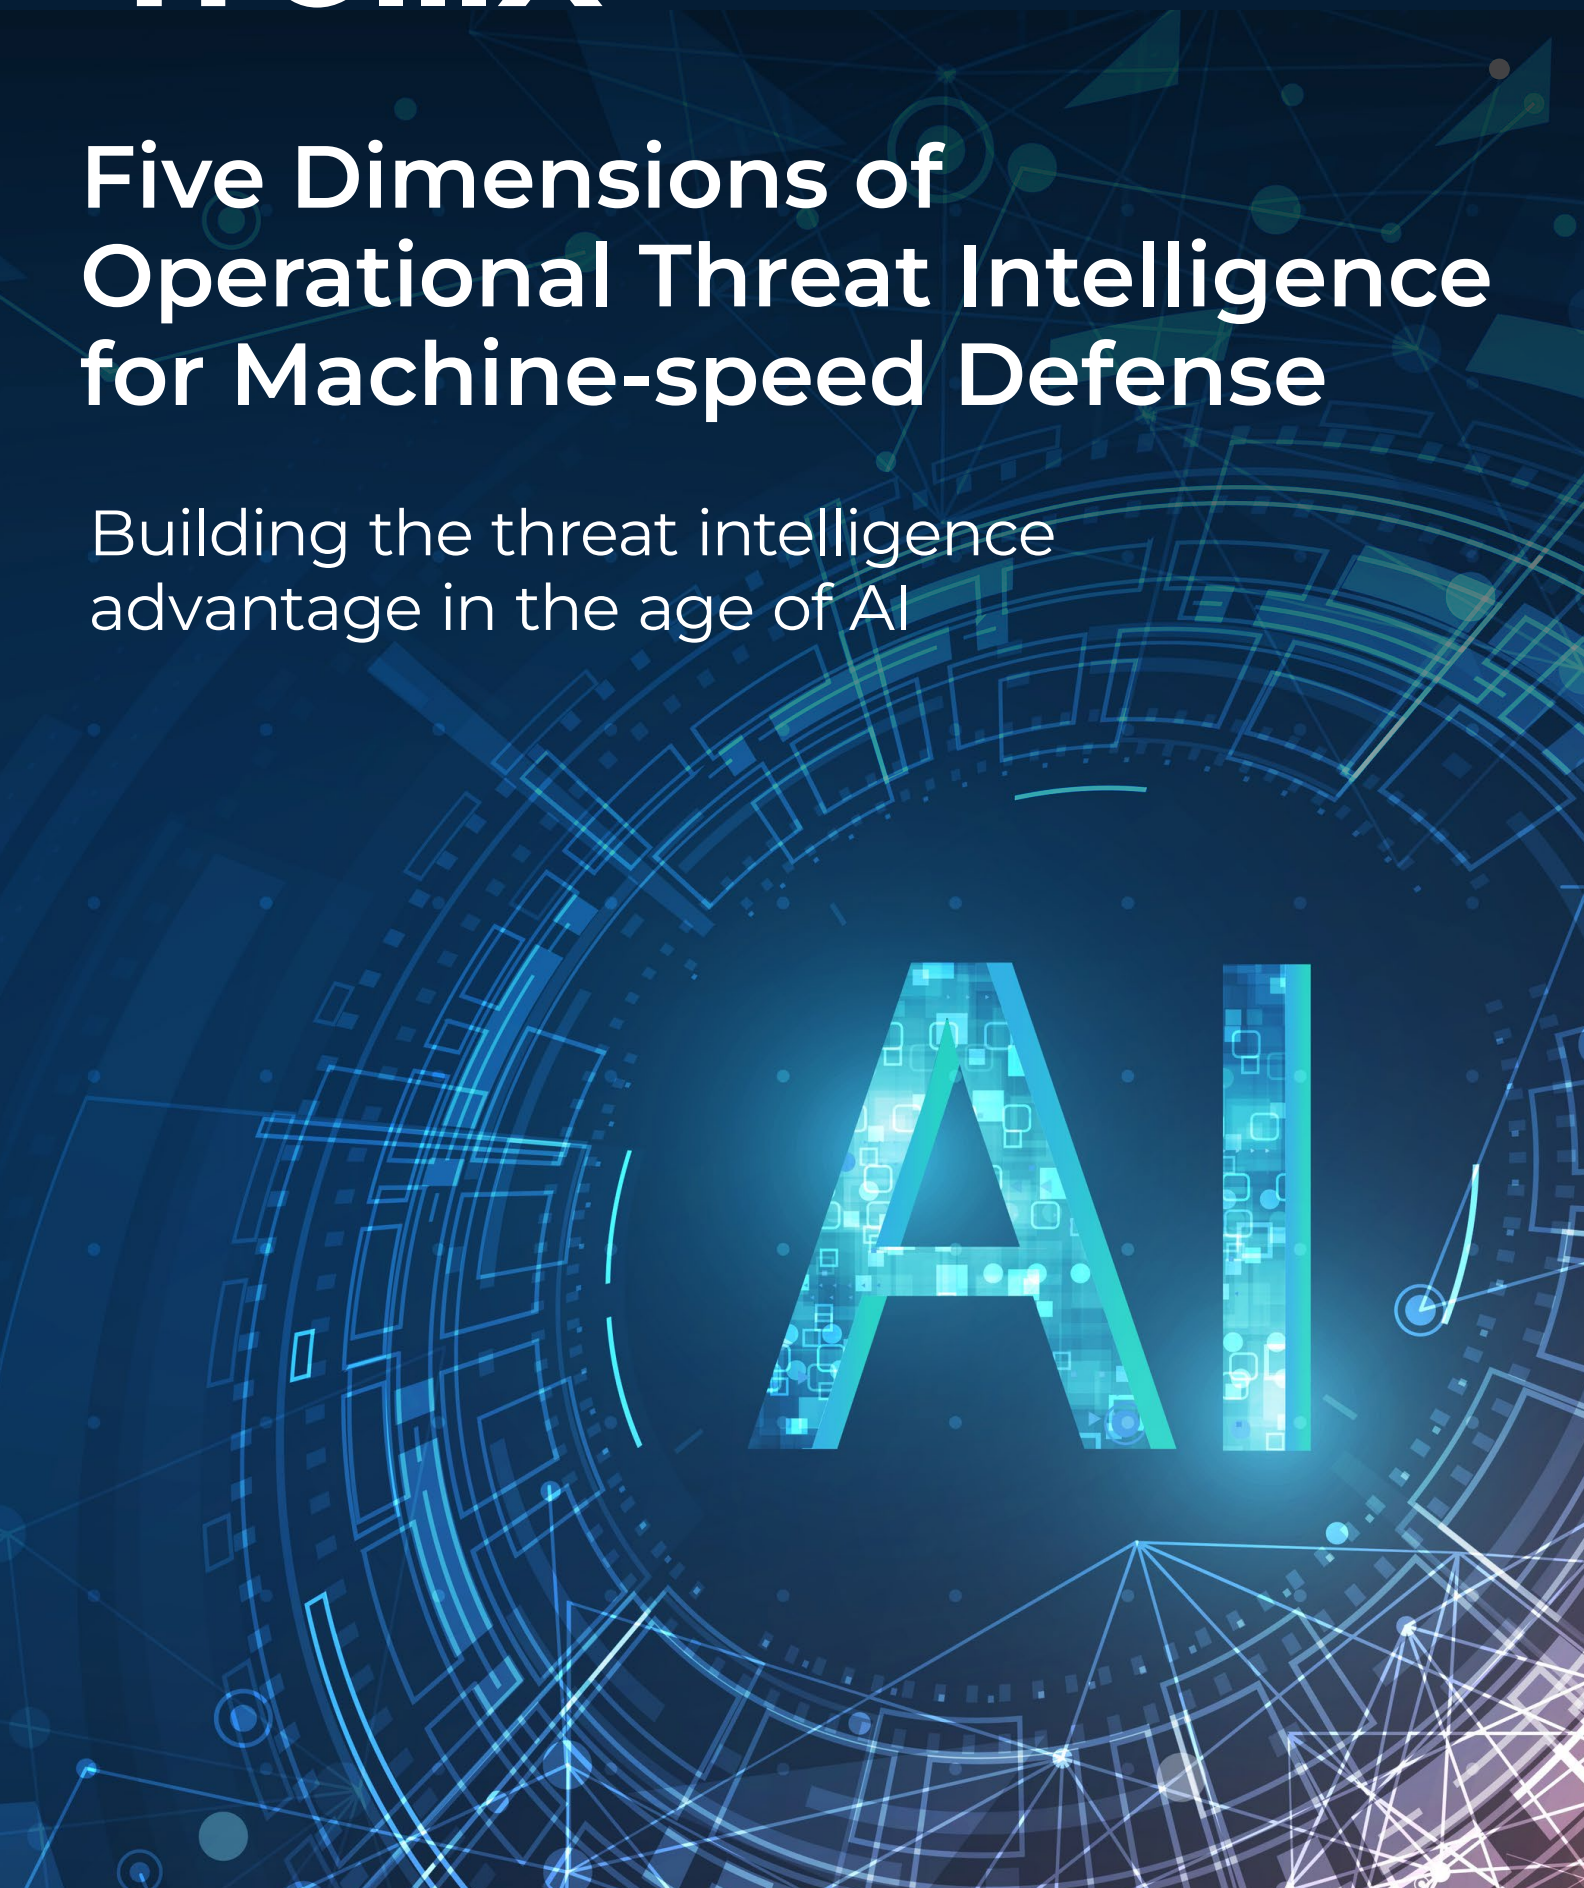

Artificial intelligence is changing the speed at which cybersecurity operates. Defenders are increasingly using AI to interpret telemetry, accelerate investigations, support threat hunting, and automate analytical work that previously depended almost entirely on human expertise. At the same time, adversaries are applying many of the same capabilities to accelerate attack development, increase variation, automate portions of the attack lifecycle, and reduce the expertise required to conduct sophisticated operations.

The consequence is not simply that cybersecurity is becoming more automated. The time available to understand and respond to threats is shrinking.

For security organizations, this creates a requirement that goes beyond machine-speed detection or machine-speed remediation. Effective defense increasingly requires **machine-speed understanding**: the ability to interpret what is happening, connect an observation to broader adversary activity, determine what it could mean for the organization, and identify what should happen next.

That understanding depends on threat intelligence.

Yet the threat intelligence industry was largely built around workflows designed for human analysts and traditional security technologies. If threat intelligence is going to become a meaningful foundation for AI-driven security, particularly in regulated, sovereign, and high-assurance environments, operational intelligence itself has to evolve.

It needs to become **structured, contextual, machine-readable, capable of guiding investigation and defensive action, and available within the security boundary where the AI is operating.**

Together, these five dimensions provide a useful framework for thinking about what AI-ready operational threat intelligence needs to become.

Part I: The challenge—security is moving toward machine speed

The pressure to evolve starts with the threat landscape.

AI is reducing the time required to analyze software, modify malicious tooling, generate variations, automate reconnaissance, and adapt attack techniques. Capabilities that once depended heavily on specialist expertise and manual effort can increasingly be accelerated or partially automated. Recent Trellix research into the weaponization and commoditization of AI in cybercrime illustrates how these capabilities are beginning to move from experimentation toward practical use within the criminal ecosystem.

AI-assisted vulnerability research introduces a similar dynamic. As systems become more capable of identifying weaknesses, analyzing exploitability, and assisting both remediation and offensive research, the traditional period between vulnerability discovery and meaningful exploitation is likely to continue compressing.

For defenders, the implication is important. Patching faster remains necessary, but patching alone does not answer whether an adversary exploited the vulnerability before the exposure was removed. Blocking a malicious process does not establish whether the attacker previously obtained credentials. Resolving a suspicious network connection does not establish whether the activity was an isolated event or part of a broader intrusion.

The fundamental security question therefore becomes larger than whether an individual alert was handled correctly. Organizations need to understand what the event represents, how it may relate to known adversary behavior, what else may have occurred, and what additional evidence should be sought.

That is the difference between responding at machine speed and understanding at machine speed.

It is also why threat intelligence becomes more important as AI adoption grows.

THE MISSING MIDDLE OF THREAT INTELLIGENCE

Threat intelligence has traditionally operated across tactical, operational, and strategic layers.

Tactical intelligence provides immediately actionable information about malicious files, infrastructure, IP addresses, domains, URLs, and other observables. It supports rapid identification, blocking, detection, and enrichment and remains an essential component of cyber defense.

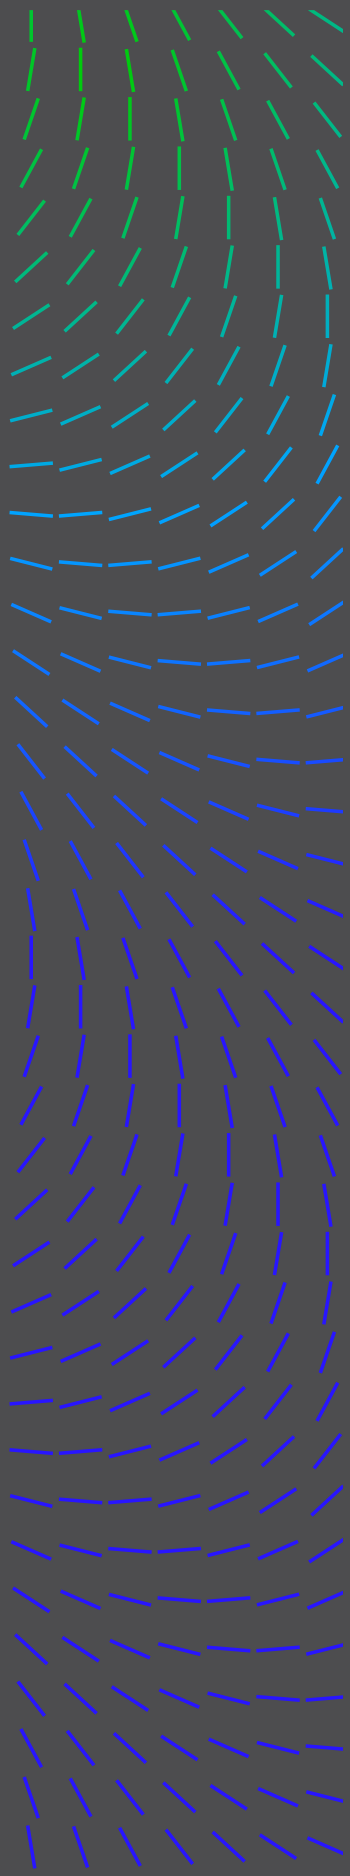

Strategic intelligence provides the broader view. It helps security leaders understand which adversaries matter, how geopolitical and technological developments may influence the threat environment, which industries and regions are being targeted, and how organizational risk is changing.

Between those layers sits operational intelligence: the knowledge of how adversaries actually conduct their operations.

Operational intelligence connects threat actors with campaigns, tools, techniques, behaviors, infrastructure, vulnerabilities, targeting patterns, and previously observed attack chains. It allows a defender to move from knowing that something is suspicious toward understanding what that activity may represent and what an adversary could attempt next.

The industry possesses enormous amounts of this knowledge. Researchers understand how campaigns evolve, malware analysts understand how tooling is employed, threat hunters recognize recurring behaviors, and incident responders understand which activities tend to appear together during an intrusion.

The difficulty has historically been translating that collective knowledge into something security systems can consistently consume and operationalize.

Much of it remains embedded in narrative reporting, threat assessments, analyst notes, research repositories, and human expertise. Converting those insights into detections, hunting logic, investigation procedures, or automated workflows still requires considerable interpretation.

In many ways, operational intelligence has become the missing middle: too complex to reduce to a simple indicator match, but too valuable to remain primarily trapped inside reports and human workflows.

AI provides an opportunity to change that, but only if operational intelligence itself is prepared for machine consumption.

Part II: The five dimensions of AI-ready operational threat intelligence

DIMENSION 1: STRUCTURED—CODIFYING WHAT WE KNOW ABOUT THE ADVERSARY

The first requirement is structure.

Threat intelligence has historically captured a great deal of adversary knowledge in prose. Research reports can explain that an adversary operates particular campaigns, uses certain tools, relies on specific techniques, targets particular industries, and tends to perform activities in recognizable sequences.

A human analyst can read those reports and mentally connect the relationships.

An AI system should not have to rediscover all of them every time it encounters a new security event.

Structured intelligence means explicitly representing the entities that make up the threat landscape and the relationships between them. Threat actors connect to campaigns, campaigns connect to tools and infrastructure, tools connect to techniques and behaviors, vulnerabilities connect to exploitation activity, and observations connect back to the broader adversary operation in which they were seen.

Structure should also preserve qualities such as confidence, provenance, supporting evidence, and temporal context. Knowing that two entities are associated is useful; knowing why the association exists, how strongly it is assessed, and when it was last observed makes the intelligence substantially more valuable.

This is the evolution from documenting adversary activity toward **codifying adversary understanding.**

Achieving this requires intelligence organizations to develop consistent data models, taxonomies, relationships, and analytical standards so that what researchers know is captured in a repeatable way. Open platforms and standards such as MISP can provide an important foundation, but the underlying discipline matters more than the individual technology.

The objective is not structure for its own sake. It is to preserve the result of human intelligence analysis so that it can be reused by analysts, products, automated workflows, and AI systems without requiring the same relationships to be reconstructed repeatedly.

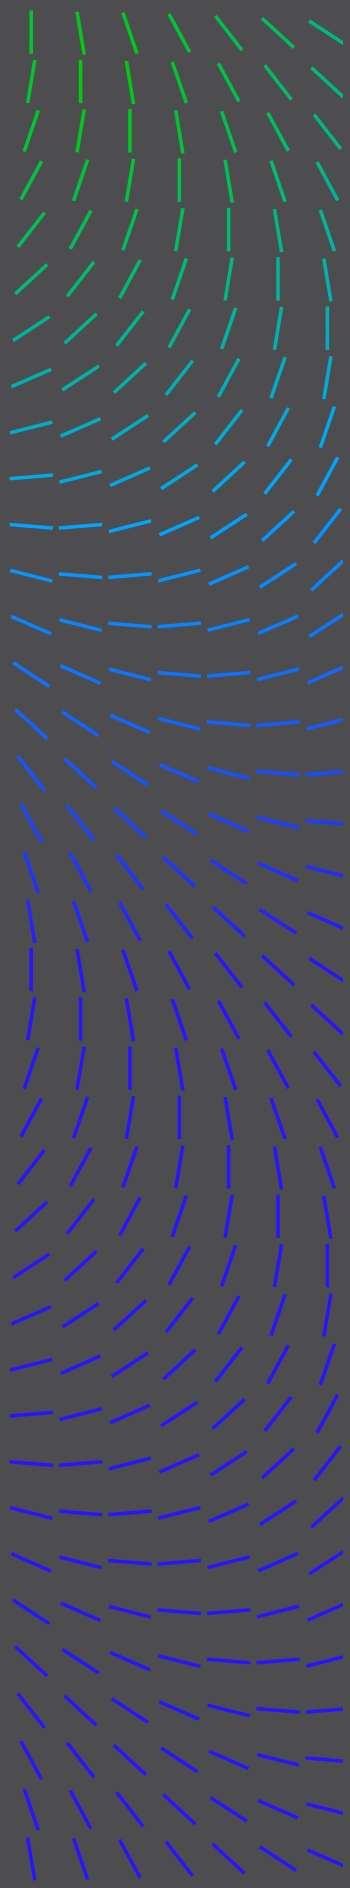

DIMENSION 2: CONTEXTUAL—GIVING BEHAVIOR MEANING

Structure explains how intelligence is represented. Context explains why an observation matters.

This distinction becomes increasingly important as adversaries gain greater ability to change the artifacts associated with an attack. Malware can be modified, infrastructure can rotate, commands can change, and individual indicators may have shorter operational lifetimes.

The underlying objectives of an intrusion are more constrained. An adversary may still need to gain access, understand the environment, obtain credentials, increase privileges, maintain persistence, move between systems, access information, or create operational impact.

Modern endpoint and network security technologies already identify many of these behaviors. The intelligence opportunity is to place those behaviors into context.

Traditional enrichment may ask whether a particular hash, domain, or IP address has been observed before. Contextual operational intelligence asks broader questions: whether the activity resembles behavior seen during previous campaigns, which adversaries have historically used comparable approaches, what other activity tended to accompany it, and how relevant those patterns are to the organization being defended.

Good contextual intelligence therefore requires more than accumulating additional data. It requires understanding the relationships between behavior, adversary intent, campaign history, targeting, industry, geography, and mission.

It must also preserve uncertainty. A behavioral similarity should not automatically become an attribution. Adversaries frequently share techniques, tools, and tradecraft. Context should strengthen investigative hypotheses rather than manufacture certainty where the evidence does not support it.

For AI, this distinction is critical. The objective is not to give a model enough information to confidently label every event. It is to provide enough context for the model to reason about what an observation may represent and determine what evidence would make that assessment stronger or weaker.

This is where behavioral intelligence can become a bridge between modern endpoint detection and response (EDR), network detection and response (NDR), endpoint, third-party telemetry, and the accumulated knowledge of threat researchers.

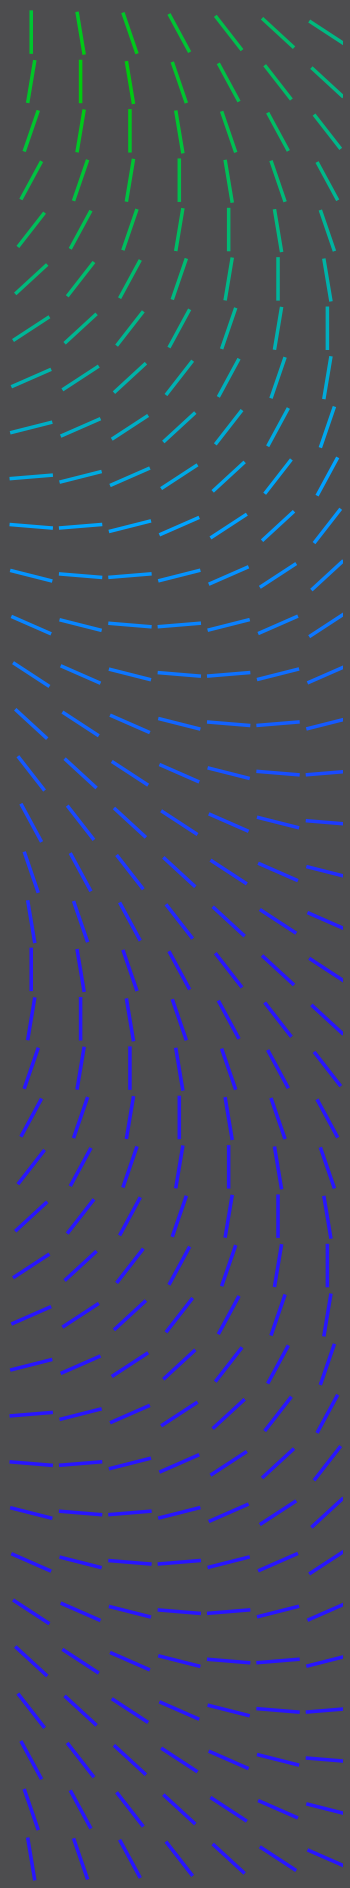

DIMENSION 3: MACHINE-READABLE—MAKING INTELLIGENCE CONSUMABLE WITHOUT ANOTHER TRANSLATION LAYER

Structured and machine-readable intelligence are related, but they are not the same thing.

Intelligence can be well organized and still depend on a human analyst to interpret it before a security product or automated workflow can use it.

Machine-readable intelligence goes one step further. It means representing and delivering intelligence consistently enough that security systems, APIs, automated workflows, and AI agents can consume it directly while preserving its meaning.

The distinction matters because the next generation of threat intelligence consumers will increasingly include machines.

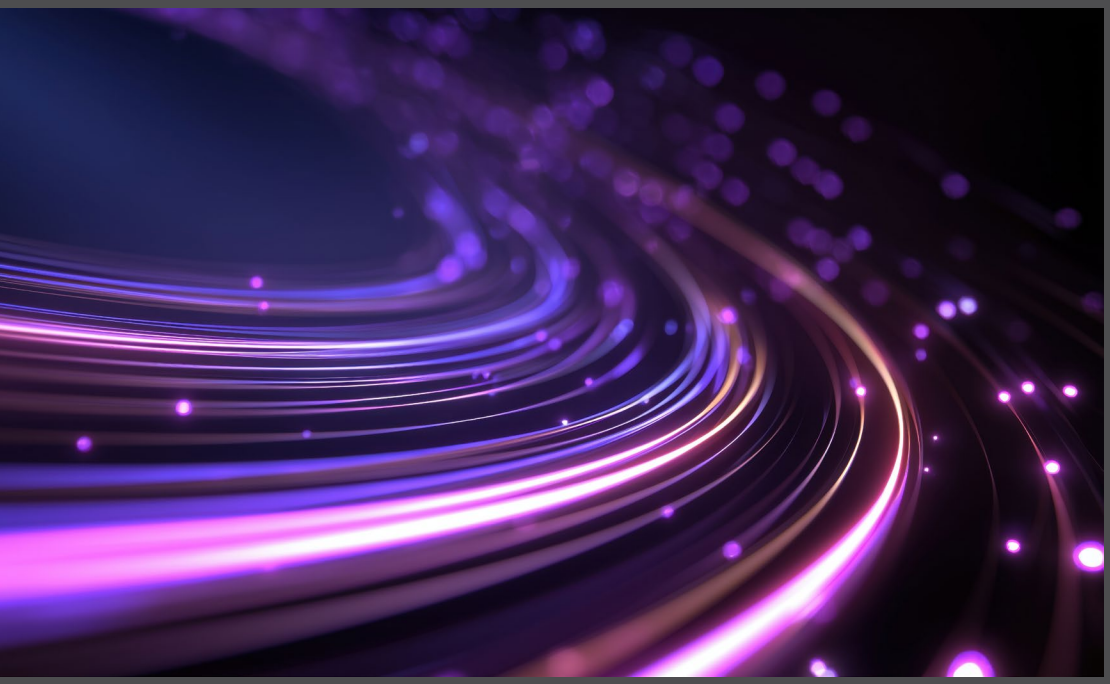

Historically, a researcher could publish a detailed report and rely on a threat hunter, detection engineer, or analyst to determine what parts of that report should be converted into operational action. In an AI-driven environment, that translation needs to become much more direct.

An AI system should be able to identify which adversaries are associated with a campaign, which techniques and behaviors characterize their operations, what confidence supports those relationships, which sectors are relevant, and what observations would be useful during an investigation without first requiring another analyst to manually extract every relationship from narrative text.

Achieving this requires consistent schemas, identifiers, taxonomies, relationships, confidence models, and methods for distributing intelligence across different consumers. It also requires intelligence producers to think about downstream machines as first-class consumers alongside human analysts.

This does not mean eliminating narrative reporting. Strategic reports, analysis, and human explanation remain essential. It means ensuring that the underlying intelligence supporting those reports can also travel independently into products, APIs, hunting workflows, and AI systems.

Threat intelligence therefore begins to serve humans and machines from the same trusted foundation.

DIMENSION 4: ACTION-ORIENTED— GUIDING WHAT HAPPENS NEXT

The fourth dimension is what separates useful machine-readable intelligence from a passive knowledge repository.

Operational intelligence needs to help guide action.

Traditional enrichment frequently answers a question about the event in front of the analyst: what is known about this file, this infrastructure, or this alert?

AI-ready operational intelligence should also help determine what should happen next.

If an observed behavior resembles patterns associated with previously analyzed campaigns, intelligence can provide context about which other behaviors have historically appeared during those operations, which systems or identities may warrant further scrutiny, and what observations would support or weaken the initial hypothesis.

This reflects how experienced threat hunters already work. They rarely begin with complete information. They develop a hypothesis, use their understanding of adversary behavior to identify what evidence should exist, search for that evidence, and continuously revise their assessment based on what they find.

AI offers the opportunity to bring elements of that methodology to a much broader set of investigations.

An alert can therefore become the starting point for an intelligence-informed investigation rather than an isolated event that simply needs to be enriched and closed. Operational intelligence can guide the questions the AI or analyst should ask next, while security telemetry provides the evidence needed to answer them.

The goal is not autonomous action for its own sake. It is greater investigative assurance.

For security leaders, the distinction is significant. Resolving one alert does not necessarily establish that an intrusion has ended. Intelligence-informed hunting can help determine whether related evidence exists elsewhere in the environment and whether the original event represents an isolated occurrence or part of a larger attack.

The same principle applies proactively. Operational intelligence should help organizations understand which adversaries are most relevant to their industry, geography, mission, and exposure and translate that understanding into questions about readiness. Do the appropriate controls exist? Is relevant telemetry available? Can the environment detect the behaviors associated with those threats? Can targeted hunts be performed before an alert appears?

At machine speed, intelligence cannot stop at explaining the threat. It needs to help guide the defensive decision that follows.

DIMENSION 5: AVAILABLE WITHIN THE SECURITY BOUNDARY—BRINGING INTELLIGENCE TO SOVEREIGN AI

The fifth dimension becomes particularly important for regulated, mission-critical, and high-assurance organizations: intelligence must be available where the AI is permitted to operate.

Much of today's AI architecture assumes connectivity. A cloud-based system may retrieve additional information, query another service, compare external sources, or request additional context whenever uncertainty arises.

Those repeated retrieval and enrichment cycles are valuable because they provide additional opportunities to clarify ambiguity and improve confidence.

Sovereign and air-gapped environments change that assumption.

Consider a naval vessel operating in a disconnected environment. Its endpoints and networks still require protection, its security controls continue to generate telemetry, and its defenders must still be able to investigate suspicious activity. The adversary does not become less sophisticated simply because the vessel cannot connect to the internet.

An AI capability aboard that vessel cannot assume that it can continuously send sensitive telemetry to an external service or retrieve additional intelligence whenever it encounters uncertainty.

The intelligence has to travel with the system.

This creates two related requirements: **availability and quality.**

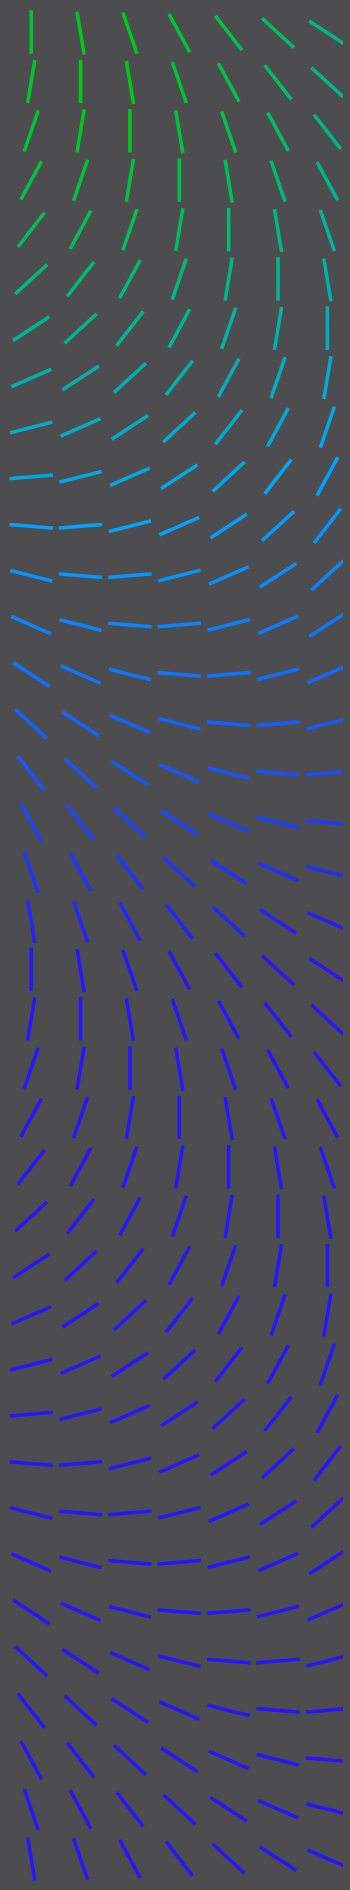

Relevant intelligence must be present within the security boundary. But simply moving a large collection of data inside the environment is insufficient. The system may have fewer opportunities to compensate for poor, incomplete, or ambiguous intelligence through another external retrieval cycle.

As a result, the intelligence available from the outset needs to meet a higher standard.

It should be curated and validated, with understood provenance and confidence. Relationships between adversaries, campaigns, behaviors, targeting, and attack patterns should be explicit rather than left entirely for the model to infer. The intelligence should also be relevant to the organization's mission and threat environment rather than representing the largest possible volume of data.

These qualities do not eliminate AI error or hallucination. They can, however, reduce the amount of unsupported interpretation the model must perform by providing a stronger foundation for its reasoning.

This is why sovereign AI should not be viewed simply as cloud AI deployed somewhere else.

Running a model locally answers where the AI executes. It does not automatically answer what the AI knows, how trustworthy that knowledge is, or how effectively it can reason when the outside world is unavailable.

In a connected architecture, AI can often reach the intelligence and iterate when more context is required. In a sovereign architecture, the intelligence must be brought to the AI, and it needs to arrive with enough structure, relevance, context, and confidence to be useful from the beginning.

For high-assurance environments, trusted threat intelligence therefore becomes part of the architecture of sovereign AI itself.

Part III: Applying the framework—from intelligence to infrastructure

These five dimensions change the way threat intelligence should be viewed in an AI-driven security architecture.

The objective is no longer simply to produce more threat data or build another enrichment service. The objective is to create a trusted intelligence foundation that can support security decisions across humans, products, automated workflows, and AI systems.

This is closely aligned with the broader Trellix Threat Intelligence mission: to turn global visibility and frontline expertise into trusted intelligence that can support human and AI-driven security decisions and turn visibility into action for **people, products, and machines.**

Trellix combines global threat visibility with security telemetry, frontline research, malware analysis, threat hunting, partnerships, and accumulated knowledge of adversary activity. That foundation spans tactical, operational, and strategic intelligence, but operational intelligence provides the critical connection between understanding a threat and acting on it.

Our opportunity is to increasingly structure and codify that operational knowledge so that it can support security systems directly.

Within Trellix Threat Intelligence, broad visibility provides the observations from which intelligence is developed. Researchers and threat hunters interpret those observations and establish why they matter. Structured intelligence preserves the relationships between adversaries, campaigns, behaviors, tools, targeting, and attack chains. Machine-readable delivery enables that knowledge to support products and AI systems, while local telemetry provides the evidence needed to determine whether those global observations are relevant inside a particular environment.

No individual component creates the advantage on its own.

Telemetry provides observation. Human expertise provides understanding. Structure preserves relationships. AI enables reasoning at scale. Security controls provide the mechanisms for defensive action.

Together, they create a trusted intelligence foundation.

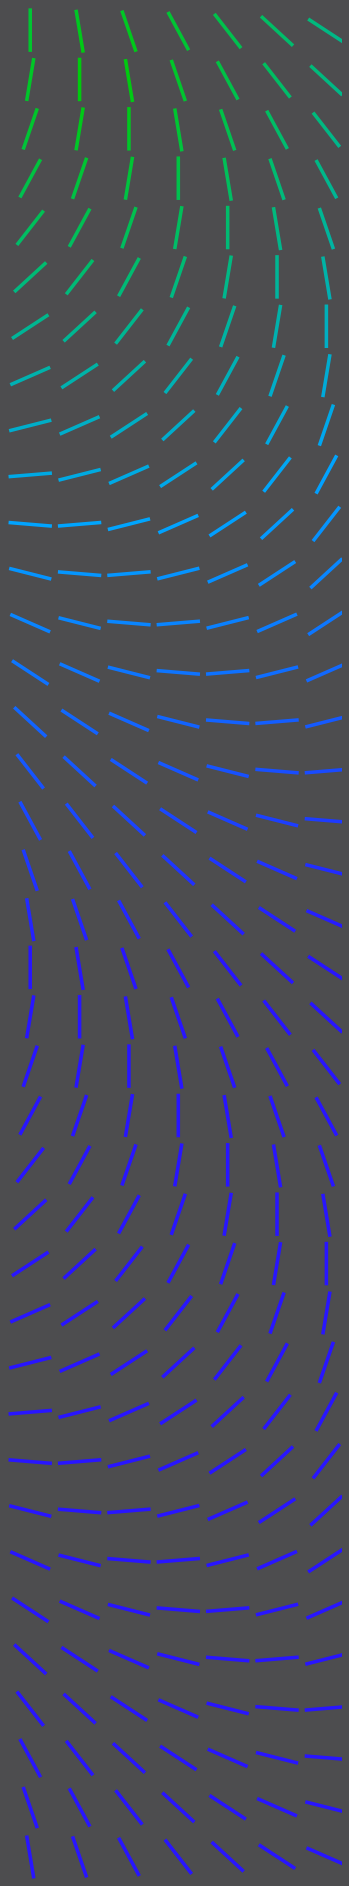

TRELLIX WISE™: APPLYING OPERATIONAL THREAT INTELLIGENCE AT MACHINE SPEED

Trellix Wise represents where these capabilities can increasingly come together.

The objective is not simply to place a general-purpose language model alongside a security operations center. It is to combine AI reasoning with trusted threat intelligence, security telemetry, and defensive capabilities so that the system can reason about both the customer's environment and the adversaries most relevant to it.

Operational intelligence provides knowledge of how those adversaries behave. Local telemetry provides evidence about what is occurring inside the customer's environment. AI can help connect the two, support analysts in interpreting activity, guide investigation, and help translate intelligence into hunting and defensive workflows.

The model therefore provides reasoning, but the value of that reasoning depends on what surrounds it.

Threat intelligence provides understanding. Telemetry provides local observation. Security controls provide action.

For sovereign and air-gapped environments, this architecture becomes particularly significant because the intelligence required to support AI-driven investigation can be provided within the organization's security boundary rather than relying on continuous access to external enrichment services.

This is where the five dimensions converge. The intelligence must be structured enough to preserve adversary knowledge, contextual enough to explain why an observation matters, machine-readable enough for AI to consume it directly, action-oriented enough to guide investigation and hunting, and available within the environment in which the AI is expected to operate.

INTELLIGENCE AS INFRASTRUCTURE

The longer-term opportunity extends beyond faster enrichment or individual AI-assisted investigations. It is the creation of a continuous relationship between global threat visibility and local defensive action.

Threat activity is observed across telemetry, research, hunting, malware analysis, partnerships, and intelligence collection. Researchers establish what those observations mean. That understanding becomes structured operational intelligence that can support humans and machines. AI applies that knowledge when interpreting local security activity and determining what evidence should be sought next. New observations can then inform intelligence, hunting logic, detections, and future defensive decisions.

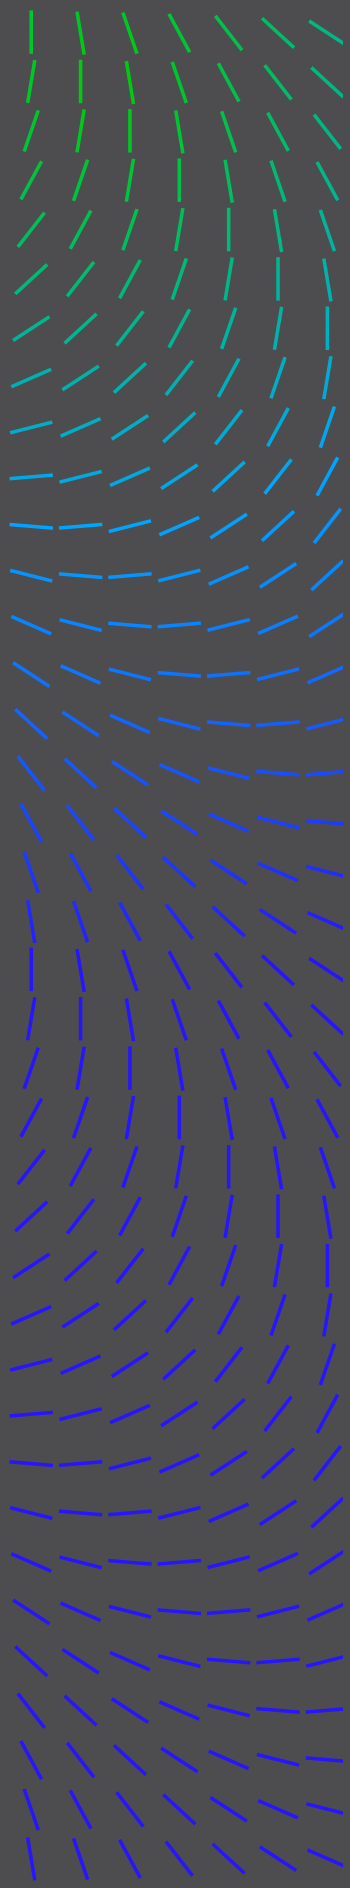

This creates a reinforcing relationship between visibility, intelligence, AI, hunting, detection, investigation, and response.

It also changes the role of threat intelligence itself.

Threat intelligence has historically been treated primarily as content: reports, feeds, indicators, assessments, and briefings. Those outputs will remain important. In an AI-driven security architecture, however, intelligence increasingly becomes part of the infrastructure through which security systems understand the adversary.

That may become one of the most important sources of differentiation in the next generation of cybersecurity.

AI models will continue to improve, agent frameworks will mature, and many capabilities that appear differentiated today will eventually become broadly accessible. The lasting advantage will increasingly depend on what surrounds the model: what it can see, what it knows, how trusted that knowledge is, how relevant it is to the mission, and how quickly new understanding of adversary activity can be translated into defensive action.

For organizations where the stakes are highest, those requirements become even more important. They do not simply need AI. They need AI grounded in intelligence they can trust, deployed within architectures they control, and capable of supporting decisions even when external connectivity is constrained or unavailable.

The future of AI-driven defense will therefore not be determined solely by how intelligent the model becomes. It will depend equally on the quality of the operational intelligence surrounding that model and our ability to make that knowledge available at the speed security now demands.

The intelligence of the model matters. The intelligence we give it may matter even more.

RELATED TRELLIX RESEARCH AND PERSPECTIVES

Weaponized AI: The Commoditization of Cybercrime
Trellix Advanced Research Center

Accelerating Trellix Engineering Transformation with AI
Trellix

Discover more cybersecurity threat insights from the Trellix Advanced Research Center, and subscribe to the ARC newsletter on LinkedIn.