Ambient.ai

DETECTION IS NOT PREVENTION: WHY SEEING A THREAT SOONER IS NOT STOPPING IT

Inside Agentic Threat Analysis and Response with Ambient Threat Detection

EXECUTIVE SUMMARY

A detection system that fires the instant a threat appears can still arrive too late. By the time a weapon is visible or a fight has started, the event is already underway, and detection has done its job at the exact moment prevention has become impossible. Seeing a threat sooner is not the same as stopping it.

Closing that gap is not a matter of faster alerts, but a matter of architecture. Prevention, in any honest sense, means recognizing the behavior that precedes an incident and acting on it before the incident happens, and that takes three things working together. The platform has to detect pre-incident precursors, not only active threats, which it does across a library of more than 150 Verified Threat Signatures. The highest-severity events have to be verified by a person before they reach the customer, which happens through the Threat Detection Analysis Center (TDAC). And response has to follow quickly, through automated, policy-driven actions that the security team defines in advance and stays in command of.

Built this way, the platform gives operators the awareness and context to resolve more than 90% of alerts in under one minute, and helps them turn detection into action through automated escalation and dispatch as events unfold rather than after the fact.

This paper covers the last two stages of the Path to Agentic Physical Security: Stage 4, Agentic Threat Analysis, where the system reasons over behavior to recognize a threat as it forms, and Stage 5, Agentic Response, where it acts. Together they are where a security operation stops documenting what happened and starts changing what happens next.

WHY DETECTION ALONE NEVER PREVENTED ANYTHING

For the narrow set of events it is traditionally built to catch, detection often works, and that much is worth being honest about. In most enterprises the analytics do fire when something happens, and an operator does see the alarm. What is far less settled is how much a system never detects in the first place, how reliably it catches what it does cover, and how little it can do with a detection once it has one.

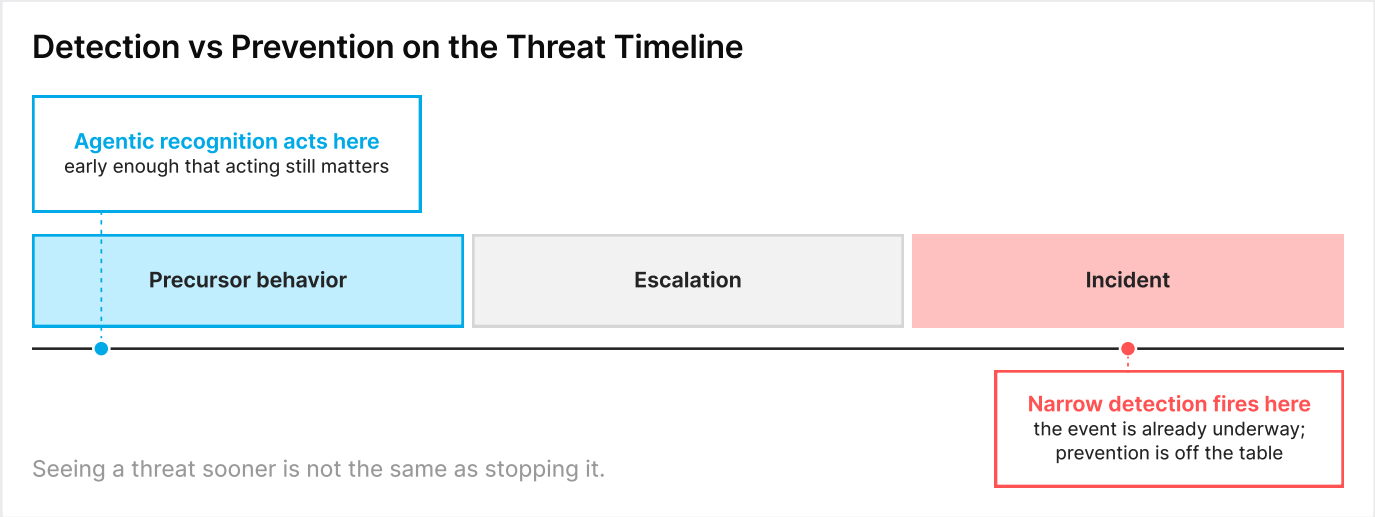

Even the best narrow detection fires at the moment of the act. Most detection is built around a single category treated as the ultimate threat, usually a weapon, and it triggers the instant that weapon appears. That instant is already the emergency. Detection has done exactly what it was built to do, at the one moment when prevention is no longer on the table.

What lets a team act early happens before the weapon ever appears, and most detection cannot see it. The warning signs are ordinary: a confrontation building, a person loitering or pacing, a hand testing a door. Reading them takes reasoning across time, and the breadth to recognize far more than a single kind of threat. Even catching a fight as it breaks out, let alone the minutes that lead to it, is beyond a detector built only to spot a drawn weapon.

Even a correct detection still has to become an action by hand. It lands in a queue. An operator assesses it, pulls up the footage, decides it is real, and works a phone tree to whoever needs to respond. Each step is human and serial, and the more urgent the event the more that delay costs, because the most serious threats are the ones least able to wait in line.

Piling on more detection without reasoning just adds noise. More detectors or higher sensitivity, on a system that cannot weigh relevance, produce more alarms rather than better ones. Operators learn to half-trust a channel that cries wolf, and the genuine event arrives in the same undifferentiated stream as the hundredth false one.

Each of these is a symptom of one condition: detection on its own, however broad or fast, documents a threat without stopping it. A system that recognizes an incident only as it happens faces a hard limit on how much harm it can prevent, no matter how much its detection improves. Call it the reactive ceiling: the point past which better detection alone yields no more prevention, because prevention takes more than detection.

The Reactive Ceiling

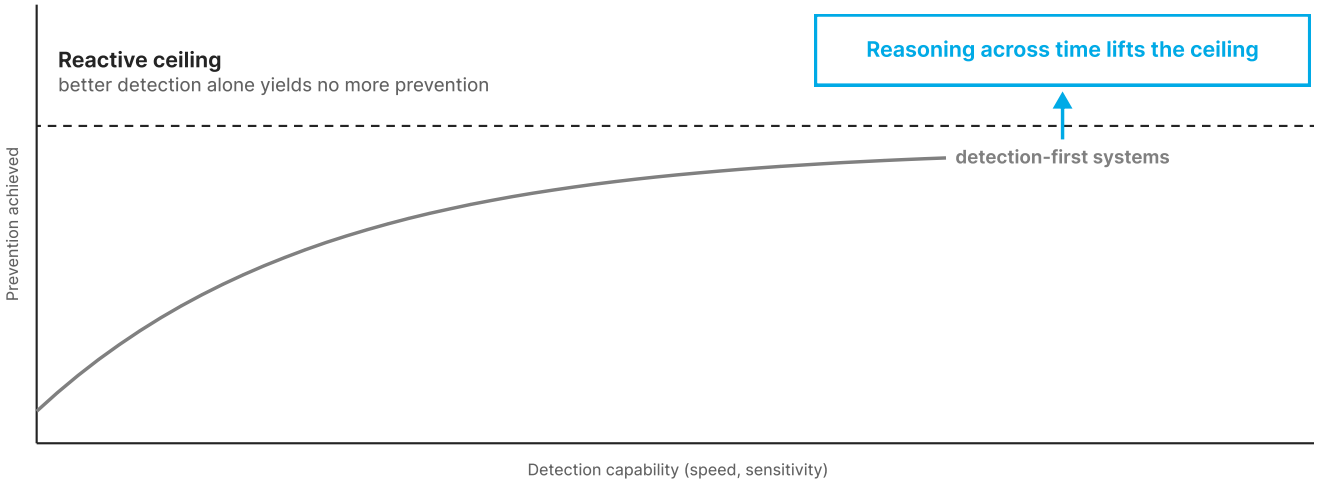

Lifting the ceiling takes a system that reasons across time, not a better-tuned detector.

THE ARCHITECTURAL GAP BETWEEN DETECTING AND PREVENTING

The reactive ceiling is not a tuning problem, and no amount of added sensitivity raises it. It is built into how detection-first systems perceive. Each of the dominant approaches improves one part of perception, and each shares the same blind spot: none reasons across time, and a precursor exists only in time.

Motion analytics (the legacy analytics built into traditional VMS) see change, not meaning. A pixel-level trigger knows that something moved, not what it was or whether it mattered. It cannot tell a delivery from a break-in, so it either floods the operator or gets tuned down until it misses the event that counted. There is no concept of a precursor anywhere in it.

Single-frame weapons detectors (point solutions such as ZeroEyes and Omnilert) do one thing, and some do it well. They scan each frame for the shape of a gun and fire when they find one. For spotting a visible weapon, that is a real and valuable capability, though not prevention. A weapon in a frame is an object, not a developing situation, and the system sees it the instant it is drawn rather than in the behavior that led to it. Detecting the gun faster is not the same as recognizing the threat sooner.

Sub-sampled cloud models (cloud-dependent platforms such as Verkada and Spot AI) miss the moments that matter most. Approaches that ship a fraction of frames to the cloud to control cost and bandwidth are, by definition, not looking most of the time. The precursors worth catching are often brief, a hand testing a door, a quick approach and retreat, and they fall between the sampled frames. An incomplete view of time cannot anticipate what happens in the gaps.

Without temporal reasoning there is no escalation awareness, and escalation awareness is the whole of prevention. Recognizing that a situation is building, and how fast, is what creates the window to act before it peaks. A system that evaluates each frame on its own never sees a situation building, only the peak itself.

The reflex fix is to put more people in front of more screens, but human review of everything is the failure mode legacy security already lives in, and it does not scale. Human judgment has a place in prevention, but only when it is aimed at the few events that warrant it rather than spread across the many that do not. Getting that division right is a problem of architecture, not staffing.

The common thread is the same in every case: you cannot prevent what your architecture recognizes only after the fact, and you cannot act on what it never recognized at all. Lifting the reactive ceiling takes a system that reasons across time and detects broadly and reliably, which is a different kind of architecture, not a better-tuned version of the same one.

Why Detection-First Approaches Hit the Ceiling

APPROACH	SEES BEHAVIOR?	REASONS ACROSS TIME?	SEES PRECURSORS?
Motion Analytics legacy VMS	○	○	○
Single-frame detectors e.g. ZeroEyes, Omnilert	◑	○	○
Sub-sampled cloud models e.g. Verkada, Spot AI	◑	○	○
Ambient Pulsar Ambient.ai	●	●	●

● Full ◑ Partial ○ None

None reasons across time, and a precursor exists only in time.

PREVENTION AS AN ARCHITECTURAL PROPERTY

Prevention is not a setting a security team turns on, but what an architecture can do when it is built to recognize a threat while the threat is still forming, early enough that acting on it still matters. Four components of Ambient Intelligence give the platform that ability. None is something an operator opens or tunes; they run continuously beneath the product, and they are why the platform can act on a threat as it forms rather than only after it has happened.

Ambient Visual Perception Engine is where pixels become description. It classifies the people, objects, and activities in each frame, producing the structured account of a scene, who is present, what they are, and what they are doing, that everything above it reasons over. Without this layer there is nothing to reason about, only raw video.

Ambient Pulsar is where reasoning happens, and reasoning across time is the precondition for anticipation. Pulsar is Ambient.ai's reasoning Vision-Language Model (VLM), and it does what detection-first systems cannot: it connects observations into a developing sequence, suppresses the routine, and elevates the behavior that signals real risk. Rather than judge a frame in isolation, it reasons over how a situation is unfolding, the way an experienced officer reads a scene, except continuously and on every stream at once. Its limit matters just as much, and the platform is built around it: on the most serious events, Pulsar narrows and prioritizes, but a human still decides.

Contextual Threat Analysis Engine turns that reasoning into an assessment. It weighs the visual, spatial, and behavioral context of what Pulsar is observing, where it is happening, who is involved, and what came before, to judge whether the activity rises to genuine risk or is merely unusual. This is the difference between flagging movement and assessing a threat.

PACS Correlation Engine adds a second kind of evidence. Shared with Ambient Access Intelligence, it fuses access events from Physical Access Control Systems (PACS) with the video of what happened at the door, and that patented correlation is the basis for every signature that depends on access data. It is what lets the platform tell an authorized entry from a tailgater following one through.

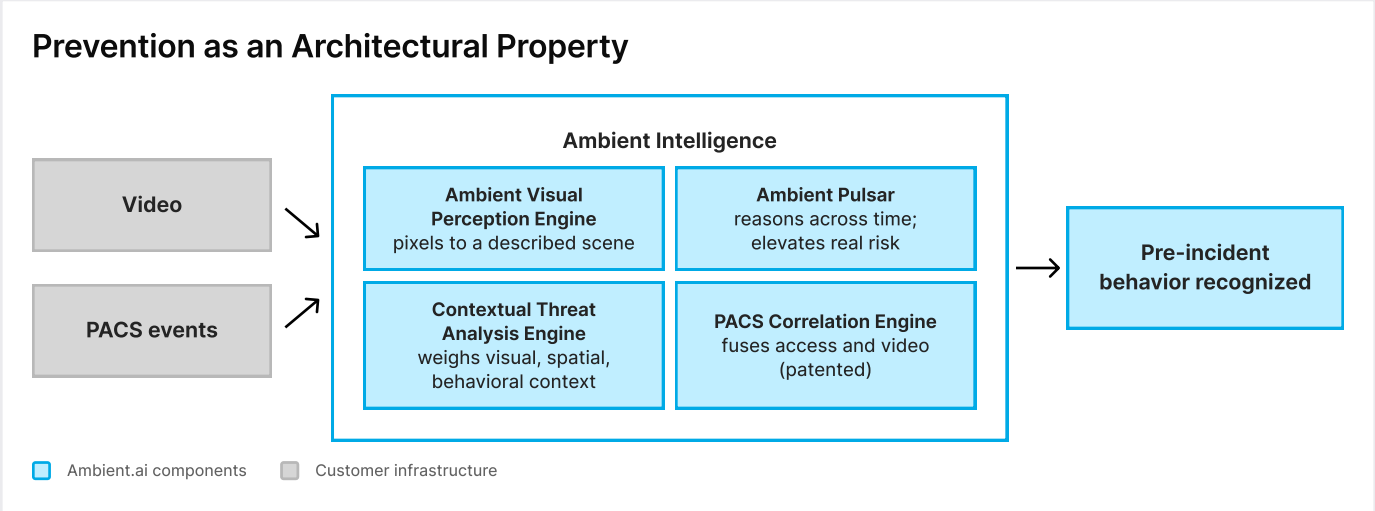

Together these give the platform something legacy detection does not have: a category for pre-incident behavior. The reasoning layer recognizes the patterns that come before an incident. Video-only loitering signatures are the clearest example: Person Loitering Outside Secure Door, Person Loitering Outside Gate, Group Loitering. None can be represented by a motion or object detector. Each is defined by behavior over time, not a single thing in a single frame.

Those signatures, the verification that protects the most serious among them, and the response they trigger are what a team actually works with.

AGENTIC THREAT ANALYSIS AND RESPONSE IN PRACTICE

An architecture that can recognize a forming threat is only worth what a team can do with it. In practice that comes down to three things: what the platform detects, which events a person verifies, and how it responds. All three rest on a single catalog of Verified Threat Signatures.

The Verified Threat Signature Library

Ambient Threat Detection reasons over a library of more than 150 Verified Threat Signatures that span the full life of a threat, from the earliest precursor to an active emergency. That breadth is part of the point. A platform that reliably recognizes a wide range of threats catches what detection built around a single category, most often weapons, never will. The selection below groups a representative set into six broad categories; the full library is available from Ambient.ai.

CATEGORY	REPRESENTATIVE VERIFIED THREAT SIGNATURES
People Safety	Person Brandishing Firearm, Person Fallen Down, People Fighting
Perimeter Intrusion	Person Jumping Fence, Person Entering Perimeter, Vehicle Loitering in Restricted Area
Suspicious Activity	Person Touching Art, Person Carrying Suitcase / Bag from Secure Room, Crowd Forming
Unauthorized Access	Tailgating, Person Entering through Propped Open Door, Invalid Badge followed by Tailgating
PACS Alarms	Fire Alarm Active with Person Exiting, Duress Button Active, Glass Break Active
Workplace Safety & Compliance	Lone Worker Presence, Animal Presence, Person Not Wearing Mask

It is important to note that PACS-related threat signatures fuse camera video and PACS events for detections no other system can make.

Verification: The Threat Detection Analysis Center

Not every detection should go to a customer untouched. The signatures with the highest stakes, Person Brandishing Firearm chief among them, are also the ones where a mistake in either direction is most costly. A missed firearm is a catastrophe, and a false one can send armed responders toward someone innocent. For this small set of high-severity signatures, a trained person verifies the detection before any alert reaches the customer.

That verification happens in the Threat Detection Analysis Center (TDAC), Ambient.ai's around-the-clock human-in-the-loop operations center. When the reasoning layer flags a high-severity event, a TDAC analyst confirms it, and only a confirmed event is passed to the customer, already validated.

The design principle is the part worth holding onto: this is human verification of the few, not human review of everything. The reasoning layer handles the overwhelming majority of detections on its own, clearing the routine so human attention is reserved for the events whose severity justifies it. That is what keeps the human in the loop without putting the human in the bottleneck. Verification adds confidence exactly where the cost of being wrong is highest, and adds no latency anywhere else.

Response: Automated and Policy-Driven

Detection and verification produce a signal a security team can trust, and Agentic Response is how the team acts on it at speed. It carries out automated, policy-driven actions chosen by threat type and severity, within rules the security team sets in advance and stays in command of. Three are generally available today.

Automated Escalation routes a verified event to the right people in the right order, so the alert reaches whoever needs it without an operator working a phone tree. **Targeted Dispatch** sends the nearest responder with the video and floor-plan context already attached, so they arrive knowing what they are walking into. **Audio Talk Down** triggers a spoken intervention through on-site speakers, which for many developing situations is enough to end them before a responder arrives.

The boundary is deliberate: the security team defines the policy and the guardrails, the platform acts within them at machine speed, and operators remain in command throughout, stepping in or taking over whenever a situation calls for judgment rather than a rule. The platform does not improvise, and it does not act outside the response the team authorized.

One Motion, from Precursor to Response

Put together, these are not three features but one continuous motion. A precursor appears, say a Person Loitering Outside Secure Door late at night. The reasoning layer recognizes it as it forms and weighs it in context instead of firing a generic alarm. If it escalates into a high-severity event, the TDAC verifies it. The moment it is confirmed, the platform acts on the team's policy, escalating to the right people, dispatching the nearest responder with context, and issuing an audio warning if that is the rule, all under the operator's command. Detection here does not end at noticing; it sets the team's chosen response in motion at once, so the right action is already underway by the time an operator turns to it, rather than waiting on a manual handoff.

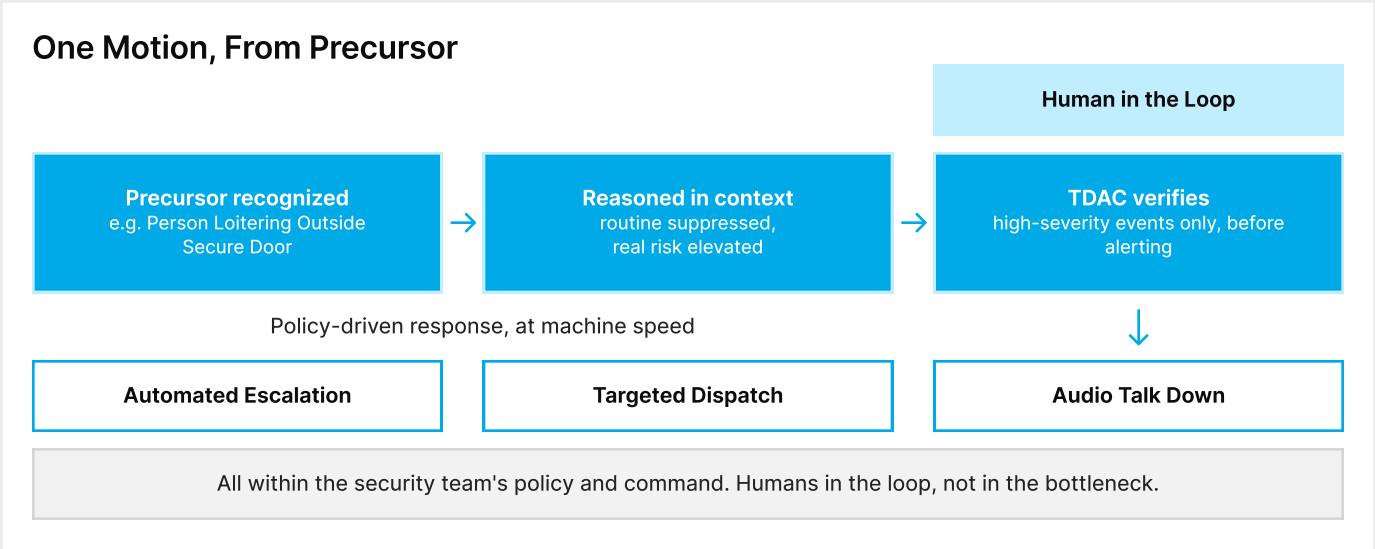

WHAT PREVENTION IS WORTH

Prevention is a strong claim, so it is worth being precise about what these numbers show and what they do not. They do not describe a world where nothing is ever attempted. They measure what changes when prevention is built into the architecture: the moment of awareness moves earlier, from after an incident to while it is still forming, and the time between awareness and action collapses. The figures below are measures of that shift.

More than 90% of alerts resolved in under one minute. When a detection arrives already reasoned, and where it matters already verified, an operator is deciding rather than investigating. With that readiness in hand, teams resolve more than nine in ten alerts in under one minute, which is the difference between a team that keeps pace with events and one that falls behind them.

10x faster response. Automated escalation and dispatch remove the manual handoffs between an alert and an action, cutting the time from detection to response by roughly an order of magnitude. The gain is not only speed. Guards and first responders move with far better information: live video, location, and context reach them as they go, so they respond to a clear, current picture rather than a secondhand description.

90% tailgating reduction. Where PACS integration is in place, the patented correlation of access events with video catches tailgating with a precision rule-based systems cannot match, and customers see it drop by roughly 90%. The outcome rests on the access-data signatures, so it applies where the doors in question are integrated.

Fewer alerts, and better ones. The quietest outcome is the most underrated. Because the reasoning layer filters out the routine and surfaces what matters, the team sees fewer alerts overall and trusts the ones it sees. Relevance is an outcome in its own right, and it is what keeps operators from learning to ignore the channel that occasionally carries the event that matters.

What Prevention is Worth

90%+ of alerts resolved in under one minute	**10x** faster response	**90%** tailgating reduction (where PACS is integrated)	**FEWER, BETTER** alerts the team actually trusts

> 71K+ behavioral threat alerts handled. 86% resolved in under one minute, boosting GSOC speed and precision.
>
> BRIAN K. TUSKAN,
> Chief Security Officer, ServiceNow

ServiceNow's security team reports the same pattern in its own operation. Across more than 71,000 behavioral threat alerts handled on the platform, its GSOC resolved 86% in under one minute. In the words of Brian K. Tuskan, the company's Chief Security Officer:

"71K+ behavioral threat alerts handled. 86% resolved in under one minute, boosting GSOC speed and precision."

None of these is a claim that incidents stop happening; each is a measure of acting earlier and faster, which is the honest content of the word prevention.

CONSISTENT DETECTION ACROSS EVERY SITE

Ambient Threat Detection runs on Ambient Foundation, the base layer of the Ambient Platform, inside the same Cloud SOC as every other module. There is no separate detection system to stand up beside the rest. The reasoning, the signatures, and the response all operate where the team already works.

Because Ambient Pulsar applies the same reasoning across every connected site, a behavior is judged by the same standard everywhere it is deployed. A precursor that matters in one building matters the same way in another, with no separate rule set per location, which is what consistent detection across a distributed enterprise actually requires. New threat concepts can be added to that shared understanding without rebuilding it from the ground up.

Two practical notes set honest expectations. Signatures that fuse access data require PACS integration: all of the PACS Alarms category, most of the Unauthorized Access signatures, and the two PACS-correlated perimeter variants. Which of them apply depends on which doors are connected. And signature selection is a deployment decision: the right set depends on what each camera can see and the environment it watches, which is work to scope at the start rather than a switch to flip.

THE COMPOUNDING SYSTEM

With detection and response, the Agentic loop closes. The platform sees across every feed, thinks across time, assesses what rises to genuine risk, and acts on it. Stage 4 and Stage 5 are where the last two steps of that loop come online.

These four capabilities, monitoring, investigations, access, and threat detection, were never four products bolted together. They are one platform reasoning over the same video in one Cloud SOC, aimed at four questions. Each makes the others better. The real-time index that makes investigations fast also feeds detection. The access verification that quiets a perimeter supplies context that detection reasons over. Detection, in turn, hands every investigation a richer place to start. The value does not add up; it compounds.

And none of it requires starting at the beginning. A team can enter the Path to Agentic Physical Security wherever its risk is sharpest, at monitoring, investigations, access, or threat detection, all on the shared Ambient Foundation, and extend from there as priorities change. The stages reinforce one another, but no stage is a prerequisite for the value of another. The principle holds across the platform: one reasoning system, adopted on the security team's terms, worth more at each stage because of the stages around it.

CONCLUSION: TURNING DETECTION INTO PREVENTION

Detection was never going to be enough by itself. Cameras have raised alarms for years, yet seeing a threat sooner has never, on its own, stopped it. The distance between detection and prevention is not a question of speed, but of what a system can recognize and how quickly it helps a team act.

Closing that distance takes three things. The platform detects the precursors that come before an incident, across a library of more than 150 Verified Threat Signatures, instead of waiting for the incident to arrive. It verifies the most serious events through the TDAC, so the alerts that carry the highest stakes also carry human judgment before they reach the customer. And it carries out automated, policy-driven response actions that the security team defines in advance and remains in command of, in real time.

The human stays in the loop, by design, on the events that warrant it. That is not a caveat to apologize for, but what responsible prevention looks like when it runs at machine speed.

Detection tells a security team what is happening; prevention means acting before it finishes, and that is the line this architecture is built to cross.