

FORENSIC DEBT: WHY INVESTIGATION SPEED IS A SECURITY OUTCOME

Inside Agentic Investigations with Ambient Advanced Forensics

EXECUTIVE SUMMARY

Most security organizations measure investigation speed as a productivity number: how many hours an analyst spends scrubbing footage, how many shifts a case consumes. That framing understates the stakes. An investigation that stays open is an unresolved question about what happened on your property, and every hour it stays open is an hour of unmanaged exposure.

This paper calls that compounding liability forensic debt. When reconstruction is slow, exposure accumulates while cases sit unresolved, evidence quality degrades as the work drags on, and senior judgment goes to retrieval instead of decisions. The debt is rarely tracked, but a security organization pays it every day.

Retiring that debt takes two layers working together. The architectural enablers inside Ambient Intelligence make meaning-based retrieval possible by indexing and reasoning over video as it streams: Ambient Real-Time Stream Indexing, the Ambient Visual Perception Engine, and Ambient Pulsar. The generally available features of Ambient Advanced Forensics then put that capability in the investigator's hands: Semantic Search, Similarity Search, License Plate Search, and Incident Timeline. Together they deliver investigations up to 20x faster, compressing days of manual review into seconds and shrinking the exposure window of every incident.

This is Stage 2 of the Path to Agentic Physical Security: Agentic Investigations, the point where reconstruction stops being after-the-fact manual labor and becomes a near-real-time tool for understanding what happened. Treated this way, investigation speed is not an operational efficiency metric. It is a security outcome.



THE ANATOMY OF A SLOW INVESTIGATION


Consider what happens in a legacy environment when an incident triggers an investigation. A package goes missing from a loading dock. A badge is used by someone who should not have it. A fight breaks out in a parking structure. The question is always the same: what happened, who was involved, and where did they go. Answering it is where the time disappears.

The timeline scrub. The work begins with one camera. An investigator pulls up the feed nearest the event and scrubs backward and forward, looking for the moment that matters. Legacy video management systems were built to record and play back, not to understand, so finding the relevant clip often means watching footage in something close to real time. One camera rarely tells the story, so the investigator moves to the next feed, then the next, reconstructing a sequence by hand across recordings that were never linked to one another.

The multi-camera tracking problem. The hardest part is following a person or a vehicle across cameras whose fields of view do not overlap. A subject leaves one frame and reappears, seconds or minutes later, somewhere else, and the investigator has to recognize them again and stitch the path together manually. Across a campus with dozens of cameras, a single subject can consume hours of review before their route is clear.

The exposure that accumulates. This is the manual failure mode legacy security inherits by design, and its cost is not only time. An open investigation is an unanswered question about a real event on a real site, and that question stays live until the reconstruction is complete. The accumulation of that exposure is forensic debt, and it compounds, because incidents do not arrive one at a time. While one case is assembled clip by clip, others queue behind it, each adding to a growing balance of unresolved exposure.

Forensic debt compounds over time



Every hour an investigation stays open is an hour of unmanaged exposure.




The labor economics. The people best equipped to interpret an incident, senior investigators, spend most of their time on retrieval rather than judgment. The scarce skill is reading intent and sequence; the work that consumes it is scrubbing and re-identifying. An organization pays for expertise and spends it on search.

The evidentiary cost. A timeline assembled slowly and by hand is more likely to contain gaps, misordered events, or missed angles. When that record has to support an audit, an insurance claim, or a legal proceeding, its quality is the difference between a defensible account and a contested one. Slow reconstruction does not just delay the answer. It degrades it.

The scale multiplier. All of this works against the security team rather than with it. Add cameras and the number of feeds to cross-reference rises. Add sites and the paths a subject can take multiply. Manual investigation time does not grow in step with the infrastructure. It grows faster, because every new camera is another feed to scrub and another handoff to reconstruct by hand. The capability meant to improve security quietly raises the cost of using it, and raises the forensic debt along with it.

The slow investigation is not a discipline problem or a staffing problem. It is what a recording architecture produces when it is asked to do the work of understanding.



WHY FASTER SEARCH IS NOT FASTER INVESTIGATION

The intuitive fix for a slow investigation is a faster search. If retrieval is the bottleneck, the thinking goes, then better indexing and better filters should solve it. The thinking is half right. Faster retrieval helps, but it does not produce a faster investigation, because retrieval and investigation are not the same task. Each common acceleration speeds up getting to footage while leaving the reasoning to the operator.

Metadata tagging finds only what was labeled in advance. A system tags footage with a fixed set of labels such as person, vehicle, and color, then lets an operator filter on them. This works for exactly the things the label set anticipated and nothing else. Real incidents are open-ended. The behavior that matters rarely fits a predefined tag: someone testing door handles, a vehicle circling a lot, a person waiting where no one waits. A fixed label set cannot surface what it was never told to look for.



FROM RETRIEVAL TO RECONSTRUCTION

If retrieval and reasoning are different tasks, then closing the gap takes an architecture that does both. Three components of Ambient Intelligence make meaning-based investigation possible. Operators never touch them directly; they run continuously in the background, and they are the reason the features in the next section work the way they do.

Keyword and filter search require you to already know what to ask. To filter for a red jacket, you have to know the subject wore a red jacket. The tool retrieves what you specify, which means it cannot help you find what you have not yet figured out. In an investigation, the unknown is the whole point.

Bolt-on analytics index objects, not behavior. They can tell you a person appeared in a frame. They cannot tell you that the person's movement across ten minutes and four cameras describes someone casing the building. Objects in isolation are not a sequence, and an investigation is fundamentally about sequence over time.

Cloud sub-sampling leaves the index incomplete. Approaches that ship frames to the cloud for analysis sample a subset to manage bandwidth and cost. That choice buys real retrieval flexibility across large archives, and it is a reasonable trade for some use cases. For investigation it is a problem, because a sampled index is an incomplete one. The events most worth finding are exactly the transient moments that fall between sampled frames: a weapon drawn for two seconds, a brief handoff, a fast exit. Because these systems also analyze footage only after it is stored, the archive becomes searchable later rather than the moment it arrives.

Every one of these methods speeds up getting to the footage. None of them does the work the footage requires. Retrieval speed without reasoning still leaves the operator to reconstruct meaning by hand, which means the forensic debt keeps accruing no matter how fast the search returns.

Faster search is not faster investigation

METHOD	SPEED RETRIEVAL?	DOES THE REASONING?	WHY IT FALLS SHORT
Metadata tagging	●	○	Finds only the labels defined in advance
Keyword and filter search	●	○	Requires you to already know what to ask
Bolt-on analytics	●	○	Indexes objects, not behavior over time
Cloud sub-sampling	●	○	Samples frames, so transient moments are missed
Ambient Advanced Forensics	●	●	Indexes as it streams and reasons over sequence

● Full ● Partial ○ None


Retrieval gets you to the footage. Investigation is the reasoning that follows.



Ambient Real-Time Stream Indexing is the prerequisite. It continuously ingests live video and tags people, objects, and activities as the footage arrives, not after it lands in storage. This single property inverts the legacy model. A legacy archive becomes searchable only after it is recorded and processed. Here the index is built in the moment, so the entire archive is queryable the instant footage exists. Every acceleration downstream depends on it, because you cannot reason quickly over what has not yet been indexed.

The **Ambient Visual Perception Engine** turns raw frames into structured signal. It detects and classifies the people, objects, and activities in a scene. This is the step that converts pixels into something a reasoning system can work with. Where a recording is just frames, the Ambient Visual Perception Engine produces a described scene: who is present, what they are, and what they are doing.

Ambient Pulsar is where reconstruction happens. Pulsar is Ambient.ai's reasoning Vision-Language Model (VLM), and its contribution to investigation is the one legacy tools lack: it connects observations into sequences and interprets them across time. It does not evaluate frames in isolation. It reasons over behavior, appearance, and context the way a human investigator reasons over a path through a building, except continuously and at machine speed.



Together these three change what a search can mean. There is no predefined label set to limit what can be found, because the system reasons about scenes in open language rather than matching fixed tags. There is no post-hoc batch step, because indexing happens as footage streams. And there is no dependence on the operator already knowing what to ask, because the system understands the scene rather than waiting to be filtered. In the terms of the Agentic loop, investigation is the assess step applied on demand over recorded time: the same reasoning that evaluates a live scene, pointed backward at the archive.

This is the architecture that stops forensic debt at its source. When meaning is indexed as footage arrives, reconstruction no longer has to be built by hand after the fact, so the exposure window starts closing the moment an incident occurs. What makes that real for the investigator is the subject of the next section: Semantic Search, Similarity Search, License Plate Search, and Incident Timeline.



AGENTIC INVESTIGATIONS IN PRACTICE

Semantic Search

Semantic Search is the primitive the whole investigation rests on. An investigator describes what they are looking for in plain language, the way they would tell a colleague, and the system retrieves the matching moments from across the entire archive. There is no tag to select and no filter syntax to learn. A request like a person in a dark jacket carrying a backpack near the loading dock after midnight returns the footage that matches, because the system reasons about the scene rather than checking it against a fixed label set.

This is open-set search, and the distinction matters. Legacy tools are closed-set: they can only find the categories someone defined in advance, so anything outside the label list is invisible. Semantic Search is open-set, which means the query is not bounded by a predefined vocabulary at all. If you can describe it, you can search for it, including the behavior no one thought to tag ahead of time. Backed by the reasoning layer, it distinguishes a person carrying a backpack from one carrying a duffel, and a vehicle waiting from one passing through. That is precisely the capability a real investigation needs, because the thing you are looking for is rarely something you anticipated.

Similarity Search

Once an investigation has a subject, the question becomes where else they appear. Similarity Search answers it. Starting from an image, either a frame selected in existing footage or a photo uploaded directly, the system finds that subject across the camera network and traces their movement from location to location.

It does this without facial recognition. Rather than identifying anyone by face, Similarity Search matches on the other attributes that describe a person in a scene: the shape and color of their clothing, their build, and their overall appearance. The investigator follows a subject across cameras based on how they look in the footage, not on who a biometric database says they are. The manual re-identification work described earlier, recognizing a subject by eye as they move between non-overlapping cameras, collapses into a single operation. A path that once took hours to assemble by hand comes back in seconds.

License Plate Search

Vehicles get the same treatment. License Plate Search locates a specific plate and follows the vehicle across cameras, without the operator scrubbing any timeline. For investigations that turn on a vehicle's route, when it arrived, where it went, and whether it returned, the answer is retrieved directly rather than reconstructed feed by feed.




Incident Timeline

The final step is assembly. Incident Timeline lets an investigator combine the relevant scenes into a single chronological narrative across multiple cameras: the moments surfaced by Semantic Search, the path from Similarity Search, and the route from License Plate Search. Instead of a stack of disconnected clips and a mental model of how they fit, the investigator builds one coherent view of the event, ordered in time and ready to support a report, a handoff, or an audit.

One Investigation, Not Four Tools

These four features are not a toolbox of separate utilities. They compose into one motion. Take the missing package from the loading dock. The investigator opens with a natural-language query in Semantic Search and finds the moment the package was taken. A person is visible. Similarity Search traces that person backward and forward across the campus, establishing where they came from and where they went. A vehicle enters the story, so License Plate Search follows it to the gate and beyond. Each result drops into an Incident Timeline, and within minutes the investigator has a single chronological account spanning a dozen cameras. From question to reconstructed sequence, there is no camera-by-camera scrubbing anywhere in between.



That is forensic debt being retired in real time. The hours that legacy reconstruction spends accruing exposure are simply not part of this workflow, so each investigation closes before its balance has a chance to grow. And the investigator's role changes with it. In the legacy model, the human is the retrieval engine, watching, recognizing, and stitching. In Agentic Investigations, the platform does the finding and the piecing together, so the investigator's work is to direct it and judge what the sequence means. The scarce human skill, reading intent and judging what a sequence means, is finally where the human time goes. The work moves from finding to thinking.



WHAT FASTER INVESTIGATION IS WORTH

The case for Agentic Investigations is not that it is faster in the abstract. It is that speed, applied to investigation, pays down forensic debt directly. Each outcome below maps to a cost named earlier in this paper.

Up to 20x faster investigations. Ambient Advanced Forensics compresses review that once took days into seconds. That speed is not a tuning setting or a faster interface laid over the same old search. It is the direct result of indexing video as it streams and reasoning over sequence, so the investigator retrieves meaning instead of hunting through footage. The architecture is the reason the number is possible.

A shorter exposure window on every incident. If an open investigation is unmanaged exposure, then closing it faster shrinks the exposure of every incident. A reconstruction that resolves in minutes rather than days turns a multi-day open question about your site into a same-shift answer. Across an enterprise's incident volume, that is the difference between forensic debt that compounds and forensic debt that is paid off as fast as it is incurred.

Senior judgment spent on judgment. When the system performs retrieval and reconstruction, investigators spend their hours on decisions rather than scrubbing. Time-to-resolution falls, not because people work faster, but because they stop doing the work a machine should do. The scarce skill is finally spent on the scarce problem.

Evidence that holds up. A timeline assembled quickly and continuously is more complete and better ordered than one stitched together by hand under time pressure. For audits, claims, and legal proceedings, that completeness is the difference between a record that stands and one that is challenged. Faster reconstruction produces a better account, not just a quicker one.

Customers describe the shift in plain terms. SentinelOne, a global technology company, runs a lean security team across sites worldwide, and forensic search is where they feel the difference daily. As Brent Kennedy, the company's Director of Security Operations, puts it:

"Forensics-wise, the ability to look for a person wearing this color top, these color pants, or carrying a box or a laptop, Ambient.ai is super fast and very accurate."

For a reader inclined to discount vendor metrics, the more durable point is the mechanism rather than the multiple. A system that indexes continuously and reasons over time will retrieve meaning faster than one that records passively and waits to be filtered. The speed is what retires forensic debt; the architecture is what produces the speed.

**UP TO 20X**

faster investigations

DAYS → MINUTESexposure window
per incident**JUDGMENT**senior time on
decisions, not scrubbing**DEFENSIBLE**evidence that holds up
for audits and claims

“ Forensics-wise, the ability to look for a person wearing this color top, these color pants, or carrying a box or a laptop, Ambient.ai is super fast and very accurate.

BRENT KENNEDY

Director of Security Operations, SentinelOne

WORKING WITH THE CAMERAS YOU ALREADY HAVE

Ambient Advanced Forensics is built to run on the cameras a security team already operates. The Ambient Platform connects to existing camera fleets through the Ambient Edge Appliance over the standard ONVIF and RTSP protocols, so the video a team already records becomes the video the platform reasons over. There is no rip-and-replace of cameras and no parallel system to run alongside the one already in place.

Because everything operates in one Cloud SOC, Agentic Investigations is added to the environment a team already uses rather than bolted on as a separate tool. The infrastructure stays; the intelligence is what is new.

None of this means a deployment runs itself. Connecting cameras, confirming coverage, and bringing the index online is real work, and it should be planned as such. What the Ambient Platform removes is the larger burden of replacing hardware or re-platforming to reach Stage 2. The starting point is the infrastructure you already have, and the value begins there: forensic debt retired on the cameras already on the wall.

FROM RECONSTRUCTION TO ANTICIPATION

Everything that makes investigation fast also points forward. The index that Ambient Real-Time Stream Indexing builds for reconstruction is the same index that supports detection in real time. Nothing about it is specific to looking backward.

That is the quiet through-line of the platform. Reconstructing what happened and recognizing what is happening are the same reasoning applied in two directions, one pointed at recorded time and the other at the live moment. An investigation is the assess step run over the past; detection is the assess step run on the present. Seen this way, the two are not separate products but the same capability aimed differently.

It also reframes forensic debt one last time. Investigation retires the debt after an incident, closing the exposure window as fast as possible. Detection works on the other side of the event, reducing the incidents that open a window in the first place. The same reasoning that pays the debt down can keep it from being incurred.



The cross-camera primitive carries the point. Following a person or a vehicle across cameras, the core of Similarity Search and License Plate Search, is the same correlation that access verification and threat analysis extend into their own domains. The capability built for forensics becomes a building block for prevention. That is where this architecture ultimately leads: from reconstructing what happened to anticipating what comes next, the work of Agentic Threat Detection and Response.

CONCLUSION: INVESTIGATION SPEED IS A SECURITY OUTCOME

Forensic debt is real, and it compounds. Every investigation that stays open is exposure left unmanaged, and in a legacy environment that debt grows with every camera added and every incident that queues behind the last. The slow investigation is not a failure of effort. It is what a recording architecture produces when it is asked to understand.

The fix is architectural, not procedural. Faster manual search does not retire the debt, because the bottleneck was never retrieval speed. It was the absence of reasoning. An archive indexed as it streams, interpreted by a reasoning layer, and surfaced through features that reconstruct sequence is a different kind of system, and it turns investigation from an after-the-fact cost into a near-real-time tool.

This is Stage 2 of the Path to Agentic Physical Security, and it stands on its own. A team can adopt Agentic Investigations for its own sake, with no obligation to progress further, and still retire the debt described here.

Investigations no longer have to take days. With reasoning and real-time indexing, the truth is available in seconds, and that speed is a security outcome.



Ambient.ai is the category creator and leader in Agentic Physical Security. At the core of its platform is Ambient Intelligence, a breakthrough engine powered by Ambient Pulsar, the first always-on, edge-optimized reasoning Vision-Language Model purpose-built for physical security. Trusted by Fortune 100 companies across campuses, data centers, and critical infrastructure.