Ambient.ai

AI CAMPUS SECURITY: WHAT HIGHER EDUCATION SHOULD EVALUATE

Six criteria for evaluating an AI campus security platform, tested against one higher education deployment.

EXECUTIVE SUMMARY

A university evaluating an AI campus security platform is not short of vendors willing to demonstrate one. What it is short of is a way to tell them apart, because every demonstration looks convincing in a conference room and the differences that matter do not appear on a feature list. They come from how a university actually operates: open by design, an adult population that is anonymous by right rather than by oversight, an estate spread across dozens of buildings, one small sworn force answering calls around the clock, and a budget set by a committee working to the academic year.

This guide sets out six criteria drawn from those conditions, one at a time. Each is written as the objection a higher education buyer raises for good reason, then turned into a question specific enough to drop into a requirements document, set out in full below. A six-row scorecard collects the asks in the form an RFP wants them. One higher education deployment, University of Northwestern - St. Paul, is then read against all six, including the point where its evidence stops. Ambient.ai's own answers are gathered in a single section near the end, so that the criteria stay the university's rather than the vendor's.

A UNIVERSITY CAMPUS IS A DIFFERENT SECURITY PROBLEM, NOT A BIGGER ONE

The estate includes a public lecture series on a Tuesday evening, a clinic that takes walk-ins, an arena on a Saturday, and residence halls occupied at three in the morning. A corporate campus controls its perimeter and knows who belongs inside it. A university does neither, because openness is the design intent rather than a gap somebody forgot to close. The population inside is adult, mobile, and anonymous by right rather than by oversight. Whatever a university buys has to hold up in an environment nobody controls.

The instinctive response to a place you cannot close is to watch more of it. Adding cameras adds coverage, and coverage is not attention. A campus running several hundred feeds still has only whoever is on shift to look at them, and those same people are dispatching calls and writing reports. According to the National Institute of Justice, operators lose 95% of their attention after 20 minutes of monitoring. What follows is the pattern most higher education teams already live with, an estate that records everything and notices very little, useful the day after an incident and inert in the hour before one.

Adding cameras adds coverage, and coverage is not attention.

WHEN ONE OFFICER ANSWERS A CALL, THE REST OF THE UNIVERSITY GOES BLIND

At University of Northwestern - St. Paul the problem had a literal shape: one patrol officer, 15 buildings, and every call for service opening a coverage gap across everything else. A campus safety director carries a 24/7 residential population, transient traffic from admissions tours, athletics, and summer conferences, and a community that expects to be treated as adults rather than as a badged workforce. Sworn campus police have real authority across those buildings, but authority is not capacity. Headcount is set by a committee working to the academic year rather than to an incident.

Bill Owen, Director of Public Safety at University of Northwestern - St. Paul, put it plainly:

> “ The moment one of my officers responded to a call, whether a medical emergency, a car lockout, or an unknown threat, patrol stopped. That was it. There was no one left to watch anything else. The rest of campus was essentially blind.
>
> BILL OWEN
> Director of Public Safety,
> University of Northwestern - St. Paul

Hiring out of that problem was not available. An officer costs $60,000 to $80,000 a year, and at a tuition-funded institution that cost lands directly on tuition.

THE PROBLEM WITH A UNIVERSITY'S CAMERA ESTATE IS THAT IT WORKS AS DESIGNED

On a conventional camera estate, coverage becomes overwhelming rather than useful, because a few hundred feeds exceed what anyone on shift can hold in mind. Investigation stays manual, which is why one residence hall incident can cost an officer an afternoon of scrubbing footage. Door alarms and legacy motion analytics generate noise in a building where propped doors are ordinary. None of it scales in a direction that helps, because the only lever the department controls is headcount.

The picture also stays fragmented across cameras, access control, dispatch, and residence life, which files related events in different places as unrelated ones. None of that is a failure of the people running the estate, because the equipment is behaving exactly as specified. Recording is precisely what it was built to do.

WHAT SHOULD A UNIVERSITY EVALUATE IN AN AI CAMPUS SECURITY PLATFORM?

A university should evaluate an AI campus security platform on six things:

- Whether it adds intelligence to the cameras and access control the campus already owns, or requires replacing them.
- Whether its handling of identity holds up in an open-access environment with adult residents and public foot traffic.
- What its threat signatures actually cover, how much of each stream it analyzes, and which detections a human checks before an officer moves.
- How much of what it detects actually reaches an officer, and what happens to everything else.
- Whether it works for a sworn patrol force with no Security Operations Center (SOC) behind it.
- Whether it has been deployed and validated in production at the number of camera streams your campus actually runs.

Each answers an objection higher education buyers raise for good reason, in a version specific enough to drop into a requirements document.

A university should not have to buy its cameras twice

A university camera estate is rarely one estate. It is what a decade of building projects, donor-funded renovations, and departmental purchases left behind: mixed manufacturers, mixed resolutions, and analog encoders still running in a basement. A platform that adds intelligence to that estate ingests the existing streams and the existing access control events, runs its models against them, and leaves the hardware in place. A platform that replaces the estate asks a university to fund cameras a second time before producing a single detection.

What gets blurred in a demonstration is the line between physical infrastructure and video management software. Cameras, access control, sensors, life-safety systems, and the network they run on are one layer that a reasoning platform can work across without touching hardware. The Video Management System (VMS) is a separate decision with two legitimate answers: keep the VMS you have and layer intelligence on top, or consolidate that software onto the platform you are evaluating. A university that just renewed its VMS will want the first; one carrying three instances from three eras may prefer the second. Neither is the compromise.

That decision is not yours alone, which is the part that catches campus safety teams out. Campus IT owns the network the cameras run on, and usually holds a settled view on bandwidth, VLAN design, where compute sits, and what data leaves the campus network before procurement has a shortlist. Ask for supported camera manufacturers and models in writing, the minimum resolution and frame rate the models require, named access control integrations rather than a general compatibility claim, and the bandwidth consumed per stream.

By the time a weapon is visible in frame the incident has already begun.

On a university campus, a privacy setting can be reversed, but an absent capability cannot

Facial recognition and behavioral reasoning are not the same technology wearing different labels. A facial recognition system builds a biometric template of a face and matches it against an enrolled database. It answers who a person is, which requires holding a store of identities to compare against. A system that reasons about behavior asks what is happening in a frame and across a sequence of frames, and it can answer that without ever knowing whose face it is looking at.

On a university campus that distinction is not a technical footnote. The population is adult, the community includes a student press and a faculty senate, and being described as running surveillance on your own students is an institutional risk that outlasts any single deployment. The objection tends to arrive late, after a design is fixed and a budget is committed. It usually arrives from people who were not in the evaluation.

The test that holds up is whether a vendor's answer is architectural or procedural. A policy stating that the identification feature is switched off can be revised by a future administrator without anyone outside the room noticing, whereas a system with no facial recognition in it cannot quietly begin performing facial recognition. Ask whether the capability exists in the product at all or is merely disabled in your configuration, what personally identifiable information is stored and for how long, whether raw video ever leaves the campus environment, and who holds authority over retention and deletion.

By the time a weapon is visible on a campus camera, the chance to prevent has passed

A notification and a threat signature are not the same object. A notification fires on a pixel change, a motion box, or a tripwire crossed, then labels whatever set it off. A threat signature is validated through layers of intelligence before it is called a threat, first analyzing the behavior and then assessing what it means in context. Ask which of the two a platform is doing, because the answer decides its false-alarm profile.

Weapons detection is critical for a university, but it is not preventative. By the time a weapon is visible in frame the incident has already begun. Prevention lives upstream, in the signatures that fire before that moment: loitering in an unusual place at an unusual hour, tailgating at a controlled door, a person falling. A platform that detects only weapons responds faster. A platform with breadth across the threat lifecycle can prevent. Ask for the full list, then ask which of those signatures are pre-incident and which only fire once an incident is underway.

Because a shove outside a residence hall is over in seconds, the next question is mechanical. Ask whether the platform analyzes every frame on every stream continuously or samples at an interval. A sampling system has a duty cycle. Whatever falls between samples is not missed by a model: it is never seen by one. Then ask how long it takes from the event in frame to the alert on an officer's device, at the stream count the campus will actually run.

A brandished weapon and a loiterer by a bike rack should not reach an officer by the same path. The right design is tiered, not uniform. High-severity signatures pass through human verification before anyone is dispatched, because a misfire at that severity means officers running and a campus locked down. Lower-severity signatures reach the officer directly, since routing everything through a person rebuilds the bottleneck the platform exists to remove. Ask which signatures a human verifies, which arrive direct, and how much time that adds.

Whatever falls between samples is not missed by a model: it is never seen by one.

Detection volume is the wrong number, and what a university never sees is the one that matters

A propped door in a residence hall at 11pm is either a fire-safety violation or a student holding it open for a friend carrying laundry. A platform that cannot tell the two apart will send both. Detection volume is therefore a misleading number to evaluate on, because the count that matters is how many things reach a human in a week and what share of those were worth the walk.

With one or two officers on shift, a system that forwards everything unusual has moved the problem rather than solved it. Within a fortnight the officer starts treating the phone as background noise, which is a failure mode no benchmark measures. The stakes climb with the severity of the detection, which is why weapons detection deserves a harder question than whether a vendor offers it: what happens when the system is wrong. On a campus, wrong means officers dispatched, a lockdown, and a community that talks about it for the rest of the semester. A false alarm at that severity is a community harm rather than an operational annoyance.

The question most evaluations skip is what the system decides not to send. Ask to test it rather than take a number for it: recall, latency, the quality of the risk assessment, and the context that arrives with an alert. Then ask how the detections that were not escalated are recorded. Whatever the vendor describes, get the escalation path written down: what triggers it, who owns it, and what the officer receives at the end.

Most universities have no SOC, and the phone in an officer's pocket is the console

Most physical security platforms are designed around a room: a SOC with staffed monitors, an operator on shift, and a supervisor who can escalate. Most universities do not have that room. Several that do staff it with student workers during term and nobody at all in August. The question worth asking is what the system does when the person who needs to act is walking a residence hall corridor at midnight with a phone in their pocket, because that is the operating condition on most campuses most of the time.
Mobile delivery then has to carry enough for a dispatch decision: a video clip rather than a still image, the camera location on a map rather than a device name only IT would recognize, and a way to acknowledge, assign, and close the incident without writing a second report that night. Ask which mobile platforms are supported, what the alert payload contains, whether the system produces the incident record your department already keeps, and how it behaves when an officer is in a basement with no signal.

Two further questions decide whether this happens on your timeline or the next year's. Ask whether the vendor sells direct or through an integrator in your region, and if through an integrator, who owns the deployment when a camera stops reporting. Then ask what a campus your size sees in its first 30 days rather than what the platform can do at maximum configuration, because committees working to an academic calendar tend to decide in the spring for a fiscal year that begins in the summer.

A demonstration on a handful of streams tells a large university nothing

A platform that performs well on 25 streams in a demonstration tells a university running several hundred cameras across dozens of buildings very little, because the constraints that break these systems appear as stream count rises. The architectural question underneath is where the compute sits when inference runs concurrently across every feed, and what has to be bought again each time the estate grows. None of that is higher education evidence, which is why the questions below still belong in your RFP.

Scale on a university estate is rarely a single number, because a university with satellite campuses, a medical center, and an athletics complex is a multi-site deployment whether anyone calls it one or not. What matters is production evidence at roughly the stream count you intend to run, not an aggregate customer count that pools deployments nothing like yours.

Ask for the largest single-customer deployment currently in production and the stream count behind it, how long it took from contract to full coverage, and what happened to detection latency between the first hundred streams and the last. Ask how many reference customers exist at your scale and in higher education specifically. Then ask what the support model looks like once you are live, meaning who answers at 2am and whether the people who deploy the system are the people who support it.

ANY VENDOR ON A UNIVERSITY'S SHORTLIST SHOULD BE ABLE TO ANSWER ALL SIX

The six criteria condense into six rows a university can lift into a requirements document or use to score vendors side by side. The middle column is the ask, phrased roughly as it would appear in an RFP. Nothing in it is specific to one platform.

REQUIREMENT	WHAT TO ASK FOR	WHAT GOOD LOOKS LIKE
Works with the cameras and access control already owned	"Provide supported camera manufacturers and models, minimum resolution and frame rate, named access control integrations, per-stream bandwidth, and confirmation that no camera or access control hardware requires replacement. State whether our existing VMS is retained and what changes if it is not."	Named models, not a general ONVIF claim. Named access control products. A bandwidth figure IT can plan against. Both VMS paths supported.
Privacy architecture that holds on an open campus	"State whether facial recognition exists in the product or is only disabled by configuration, what personally identifiable information is stored and for how long, where raw video resides, who controls retention and deletion, and whether all access to footage is logged."	Absence of the capability rather than a setting. Raw video under institutional control. Retention authority held by the university. Readable access logs.
Threat signature depth, breadth, and reliability	"Provide the full threat signature list, marking which signatures are pre-incident and which fire only once an incident is underway. State whether every frame of every stream is analyzed continuously or sampled at an interval, the elapsed time from event to alert on a responder's device at our stream count, and which signatures a human verifies before dispatch."	A signature list, not a signature count. Pre-incident coverage well beyond weapons. Continuous analysis stated mechanically. A latency figure at our scale.
Alert fidelity and a defined escalation path	"Support an evaluation of alert fidelity by testing rather than by quoted figures, across recall rates, latency, quality of risk assessment, and the context information delivered with an alert. State the escalation path for high-severity detections and how non-escalated detections are recorded."	Fidelity demonstrated under test on all four dimensions rather than asserted. A searchable record of what was not sent.
Field operation without a security operations center	"Describe the mobile alert payload, the acknowledge, assign and close workflow, incident record output, behavior on degraded connectivity, what a campus of our size sees in its first 30 days, and whether deployment runs direct or through an integrator."	A clip plus mapped camera location. A workflow that produces the incident record already kept. A named party accountable for deployment.
Enterprise maturity and validated deployment at scale	"State the largest single-customer deployment in production and its stream count, the timeline from contract to full coverage at that scale, how detection latency behaves as streams are added, references at our scale and in higher education, and the support model with response commitments."	Production stream counts at or above your target, with contactable references. A timeline stated in weeks. A named support path.

> “ We didn’t get there by hiring more people. We got there by getting smarter. And that’s what Ambient made possible.
>
> BILL OWEN
> Director of Public Safety,
> University of Northwestern - St. Paul

UNIVERSITY OF NORTHWESTERN TESTED THESE CRITERIA ON EXISTING CAMERAS AND ONE OFFICER

University of Northwestern - St. Paul ran the estate described above: 135 fixed cameras that recorded and offered no intelligence, one patrol officer, 15 buildings. Ambient.ai deployed against that existing camera infrastructure with no hardware replacement, which is the first criterion in plain form. What changed was not coverage but what it produced: more than 100 incidents a week the team previously had no way to know about, most of them ordinary, including propped doors, loitering, and vehicles in restricted areas after hours. That is the fourth criterion, and the routine detections changed daily operations more than the dramatic ones. Response time to actionable alerts now runs under 3 minutes, reaching the patrol officer's phone rather than a monitoring desk the university does not have, which is the fifth. On signature coverage, the third criterion, Owen is specific:

> “ On any given camera stream, we are running 6 to 7 threat signatures simultaneously. That is a completely different level of coverage from our existing infrastructure without adding a single camera.
>
> BILL OWEN
> Director of Public Safety,
> University of Northwestern - St. Paul

The weapon-brandishing signature triggered within 6 to 10 hours of deployment, on a capability the team had treated as insurance. Owen's summary matters most to a campus weighing headcount against technology:

> “ We didn't get there by hiring more people. We got there by getting smarter. And that's what Ambient made possible.
>
> BILL OWEN
> Director of Public Safety,
> University of Northwestern - St. Paul

At 135 cameras University of Northwestern is a production deployment rather than a pilot, but it is not evidence about behavior at two thousand streams. A university several times that size should hold the sixth criterion to its own numbers and ask for references at its own scale.

The full account of the deployment is here: how University of Northwestern rebuilt campus safety without adding headcount.

HOW AMBIENT.AI ADDRESSES THESE SIX CRITERIA

The same six criteria, now with Ambient.ai's own answers, in criterion order. Where an answer is directional rather than numeric, that is deliberate: a figure quoted away from a specific campus rarely survives contact with one.

1. Existing cameras stay in place, and so does the capital cycle

Ambient.ai adds intelligence to a campus's existing cameras rather than requiring their replacement: it ingests live video from cameras already installed and correlates those streams with events from the access control platform already in place. Integrating reasoning into existing cameras and access control preserves prior investments and avoids downtime, which is the difference between a platform a university can run this academic year and one that waits on a capital cycle.

2. Privacy by design means no facial recognition and nothing that identifies who is present

Ambient.ai is built on privacy by design: its vision-language models (VLMs) reason about what is happening rather than identifying who is present, with no facial recognition and no personally identifiable information used for detection.

3. Signatures fire before and during an incident, and the highest-severity ones get a human check first

This criterion runs from what fires long before an incident to what fires during one, and the answers cover both ends of it, gun detection included.

Ambient.ai's models read sequences of behavior rather than single frames, which separates a delivery from a break-in and surfaces pre-incident behavior before it escalates. Ambient.ai carries more than 150 verified threat signatures spanning the full threat lifecycle, from a person loitering outside a gate through to a brandished firearm.

Ambient.ai's models run at multiple frames per second on all cameras at all times and do not rely on sub-sampling. Ambient.ai's alert latency is near real time, with detections happening in a matter of seconds.

Ambient.ai routes high-severity detections, including a person brandishing a firearm, through human-in-the-loop verification in a 24/7 operations center before the customer is alerted. Weapons detection is the worked example of that path, and the verification step adds a few seconds.

4. Confidence scoring, confidence thresholds and multi-signal gating are all system capabilities

On Ambient.ai, confidence scoring, confidence thresholds, and multi-signal gating are all system capabilities.

5. Built to run without a security operations center behind it

Ambient.ai is built to operate without a security operations center behind it, not around the absence of one. Alert escalation is automated and mobile teams receive visual context with the alert, which means a guard sees the video before responding rather than after arriving. Escalation and dispatch are initiated on the basis of threat type and severity, which is what routes an alert, triggers a dispatch, or opens audio talk down in real time without an operator deciding each step by hand. Where a university does have a monitoring desk, staff there track activity and resolution in real time while the responder is moving. Where it does not, the same design holds, because during low-traffic hours a single patrol guard can operate as both responder and SOC, reducing overhead without sacrificing coverage.

Ambient.ai supports both iOS and Android. Detections run at the edge, so a loss of network connectivity at the site does not interrupt detection. The officer receives a video clip with a full explanation of the detection, and can collaborate with remote operators to resolve the threat. Integration with third-party systems runs through APIs or webhooks.

The officer's own signal is a separate question from the site's. Ambient.ai runs on two surfaces: a cloud-based view for whoever is coordinating the response, and a phone app for the guard in the field. If the guard's phone has no connectivity, alerts continue to reach the cloud view, and the person there can keep coordinating through other communication devices or by contacting another operator. That view opens on a laptop and needs no operations center behind it. Where a campus does run a monitoring desk, it is the same view on a larger screen.

Ambient.ai sells both direct and through integrators, with specialist customer success managers supporting implementation.

6. Built to scale to thousands of feeds without added infrastructure

Ambient.ai's architecture scales to thousands of video feeds across any number of sites without adding servers, storage arrays, or infrastructure complexity, and it is trusted by Fortune 100 enterprises across corporate campuses, data centers, and critical infrastructure.

Deployment runs in a matter of days from start to finish, and a lot of that depends on customer requirements. Ambient.ai provides low latency, high-fidelity alerts at scale.

THE EVALUATION CRITERIA COME FROM THE UNIVERSITY, NOT FROM THE VENDOR

- Ask whether the platform adds intelligence to cameras and access control you already own, and settle the VMS question separately. Both VMS paths are legitimate.
- Ask whether the privacy answer is architectural or procedural. A configuration setting can be changed by a future administrator; an absent capability cannot.

- Ask which threat signatures are pre-incident and which only fire once an incident is underway, whether every frame is analyzed or sampled, and which detections a human verifies.
- Ask what the system decides not to send you, and what record it keeps of those detections.
- Ask what happens when the responder is an officer on patrol with a phone rather than an operator at a desk, because that is the normal case on a campus.
- Ask for production evidence at your own stream count, with a deployment timeline, latency behavior as streams grow, and references at comparable scale.

FOR A HIGHER EDUCATION SECURITY LEADER, THE USEFUL NEXT MOVE IS NOT TO SCHEDULE ANYTHING

None of this resolves into a shortlist, because the criteria are not a ranking of vendors. They are what a university's own operating conditions demand once somebody writes them down, which is why a demonstration booked before that happens tends to set the terms of the evaluation for you. Take the six questions into your own estate first, then read one higher education deployment as a worked example of what they surface. The list was never the vendor's to write.

Once the six are written down in your own buildings, your own camera count, and your own staffing, a conversation about a specific campus is worth having, and we are glad to have that one whenever you are ready for it.

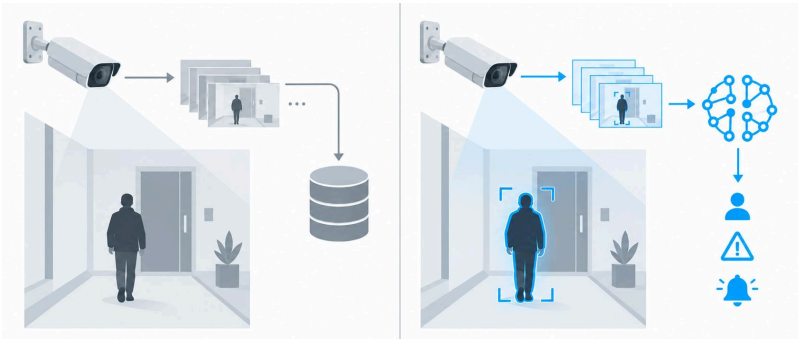

Diagram contrasting a recording-only camera estate with an intelligence layer that reasons about campus activity.

University of Northwestern's deployment, read against these six criteria