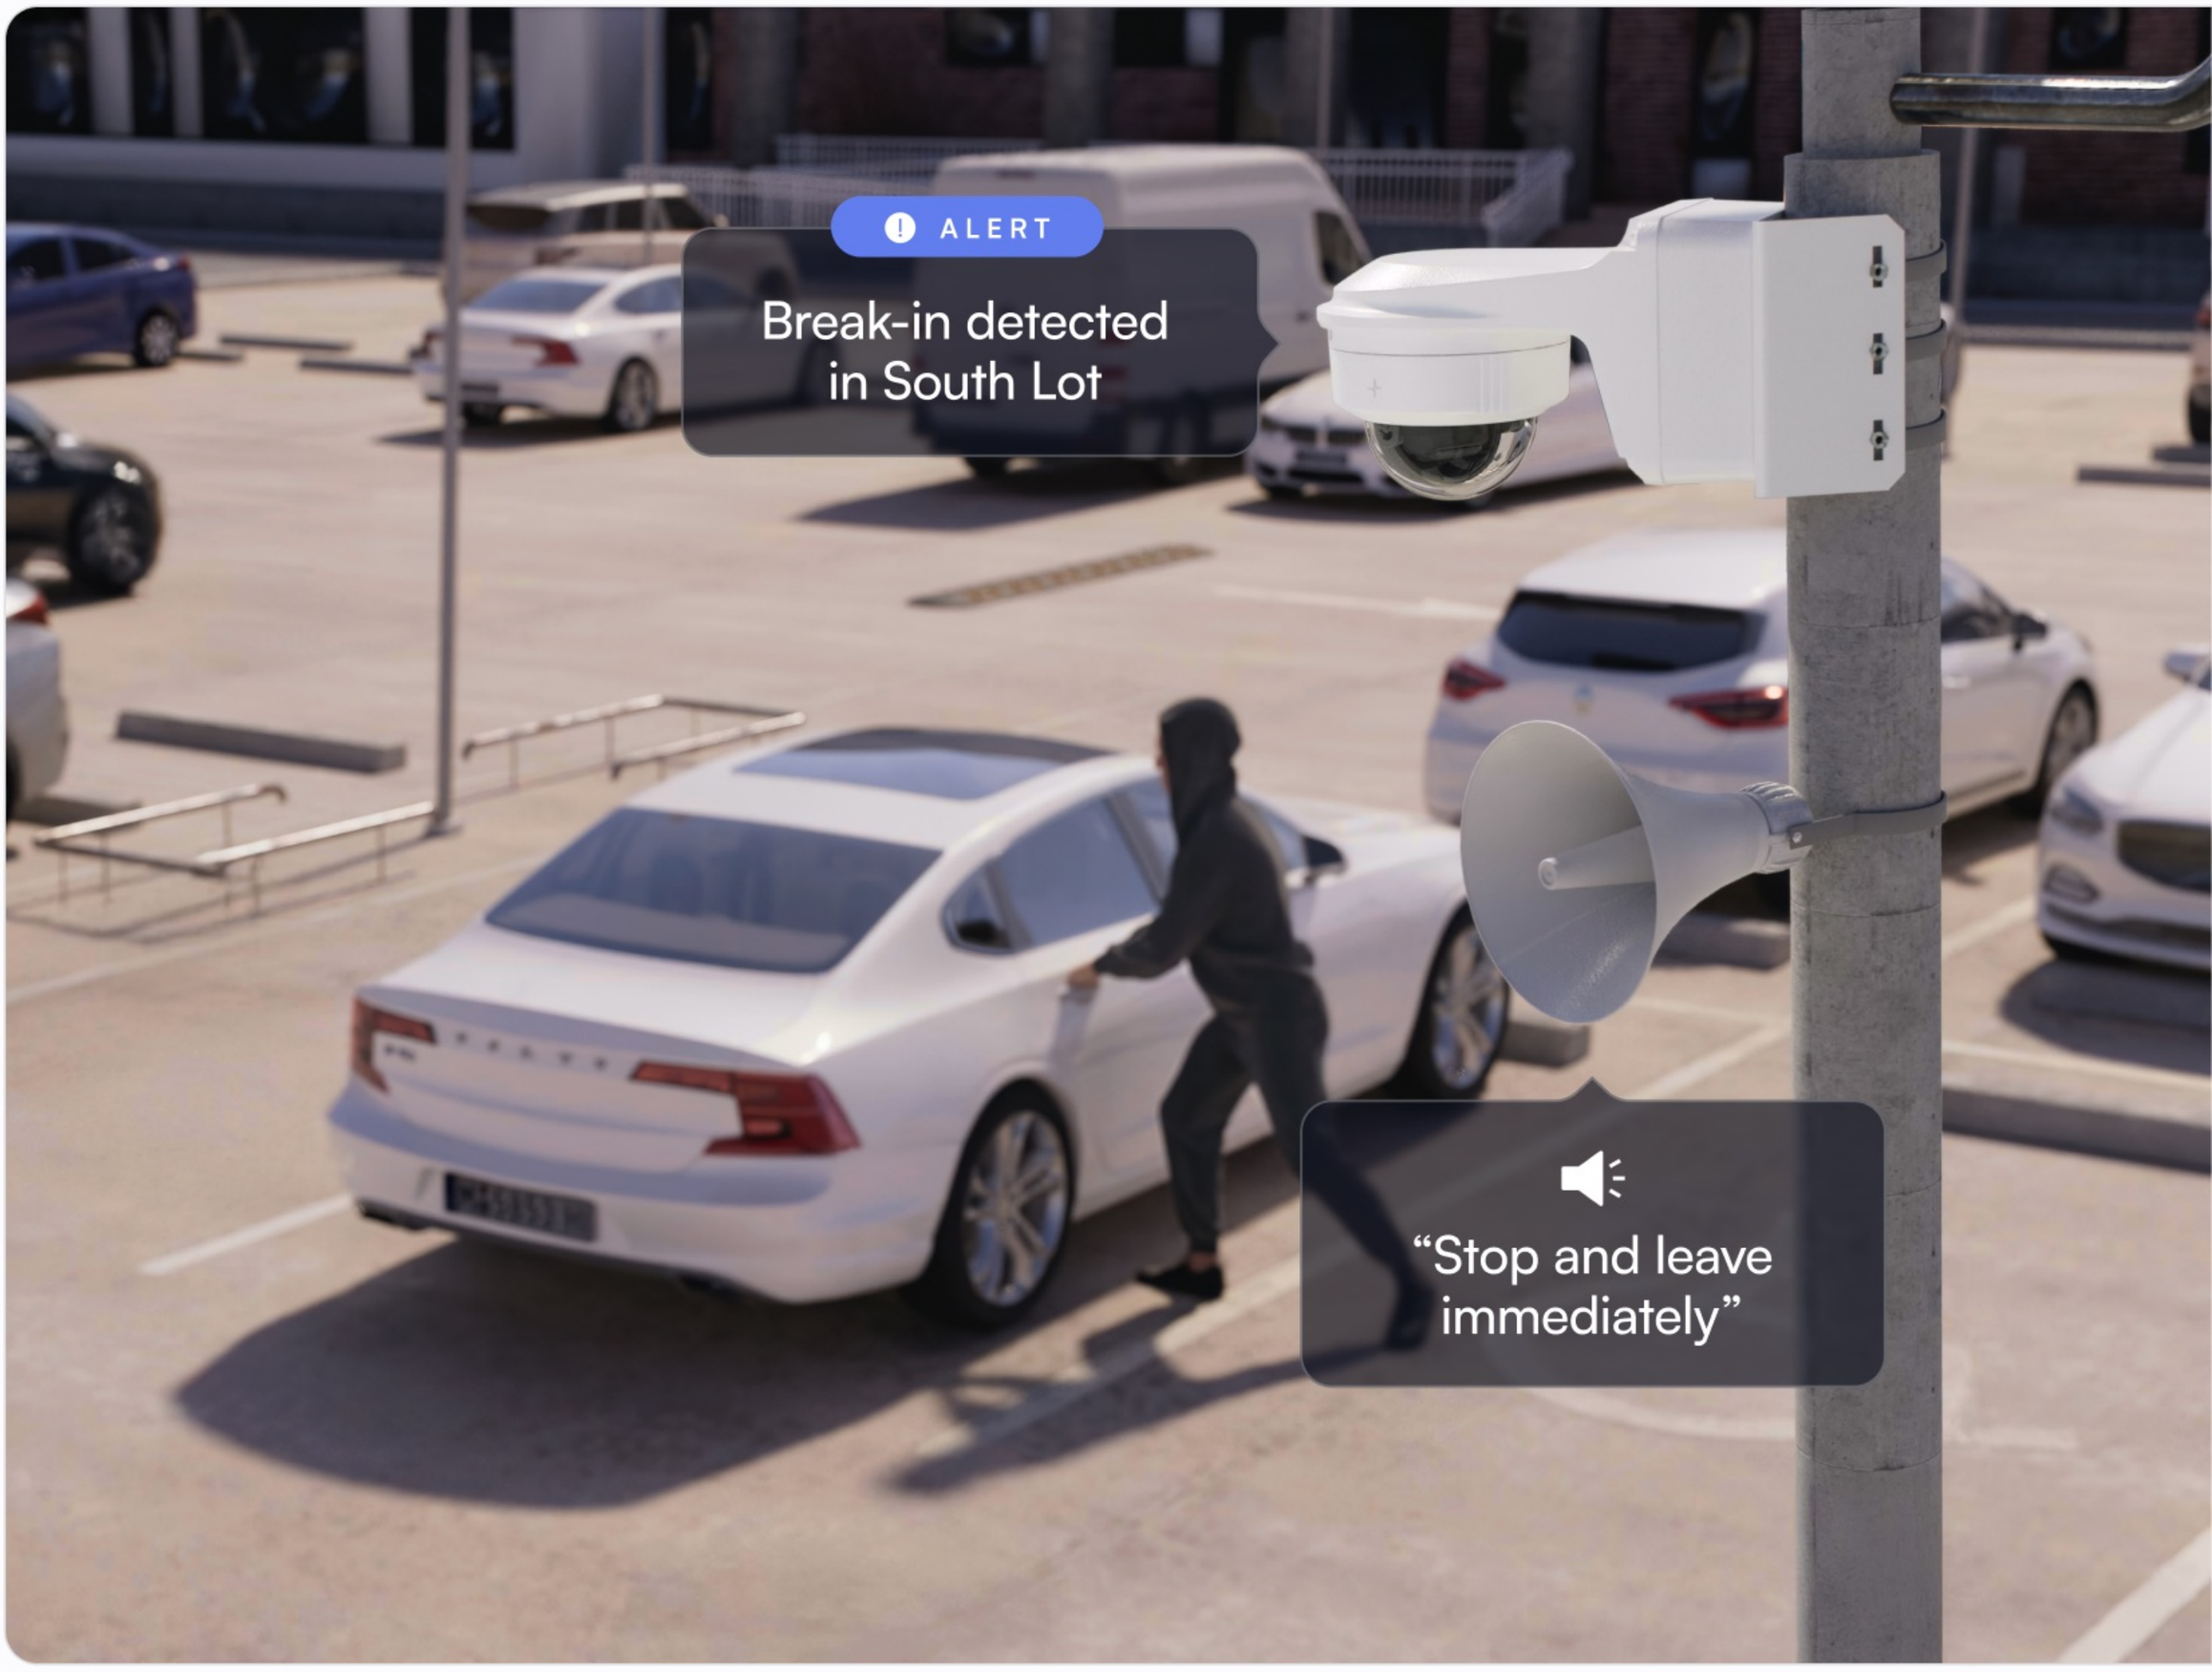

LUMANA’S DEFINITIVE BUYER’S GUIDE, 2025-2026

Evaluating AI Video Security

10 critical factors that separate truly intelligent video security from legacy systems and basic cloud cameras — and why it matters for your organization's security, safety, and operations.

What's inside

AI video security by the numbers

30.6% CAGR
AI video surveillance market through 2030
Source: Grand View Research

90%
Reduction in false alarms with AI-powered systems
Source: VORTEX

$28.8B
Projected AI video surveillance market by 2030
Source: Grand View Research

The three generations of video security

Generation 1

Legacy / On-Premise

- NVR/DVR-centric architecture
- Passive recording only
- Manual footage review
- Siloed, site-by-site systems
- Physical device export (USB)
- On-site presence required
- High hardware & maintenance const
- No operational data insight

Generation 2

Cloud Video Surveillance

- Remote access from anywhere
- Centralized multi-site management
- Scalable storage
- Basic motion-triggered alerts
- Reduced infrastructure investment
- Still requires manual monitoring
- Limited analytics intelligence
- Reactive, not proactive

Generation 3

AI Video Security

- Context-aware real-time detection
- Continuously learning models
- Proactive alerting - not just recording
- Natural-language event search
- Operations & safety intelligence
- Automated investigations
- Up to 90% fewer false alarms
- Camera-agnostic, enterprise-scale

INTRODUCTION

Security cameras are everywhere. Intelligent ones are rare.

Most organizations today operate under a false assumption: that having cameras means having security. In reality, the average security team is drowning in footage they can't act on, alerts they've learned to ignore, and systems that require them to be in the right place to respond to an incident that has already happened.

The shift from legacy video systems to cloud-managed surveillance was the first wave of modernization. The second — and far more transformative — wave is AI. Not AI as a marketing tag, but AI as the foundational architecture: systems that can perceive context, learn from their environment, and surface the right information at the right moment without a human having to watch every screen.

> *"Surveillance is no longer about passive monitoring. It's about real-time intelligence that unlocks operational efficiencies while providing businesses a proactive security edge."*

This guide gives security directors, IT leaders, and operations teams the language and frameworks to evaluate the real differences between three distinct generations of video security — and make the investment that will serve their organization for years to come.

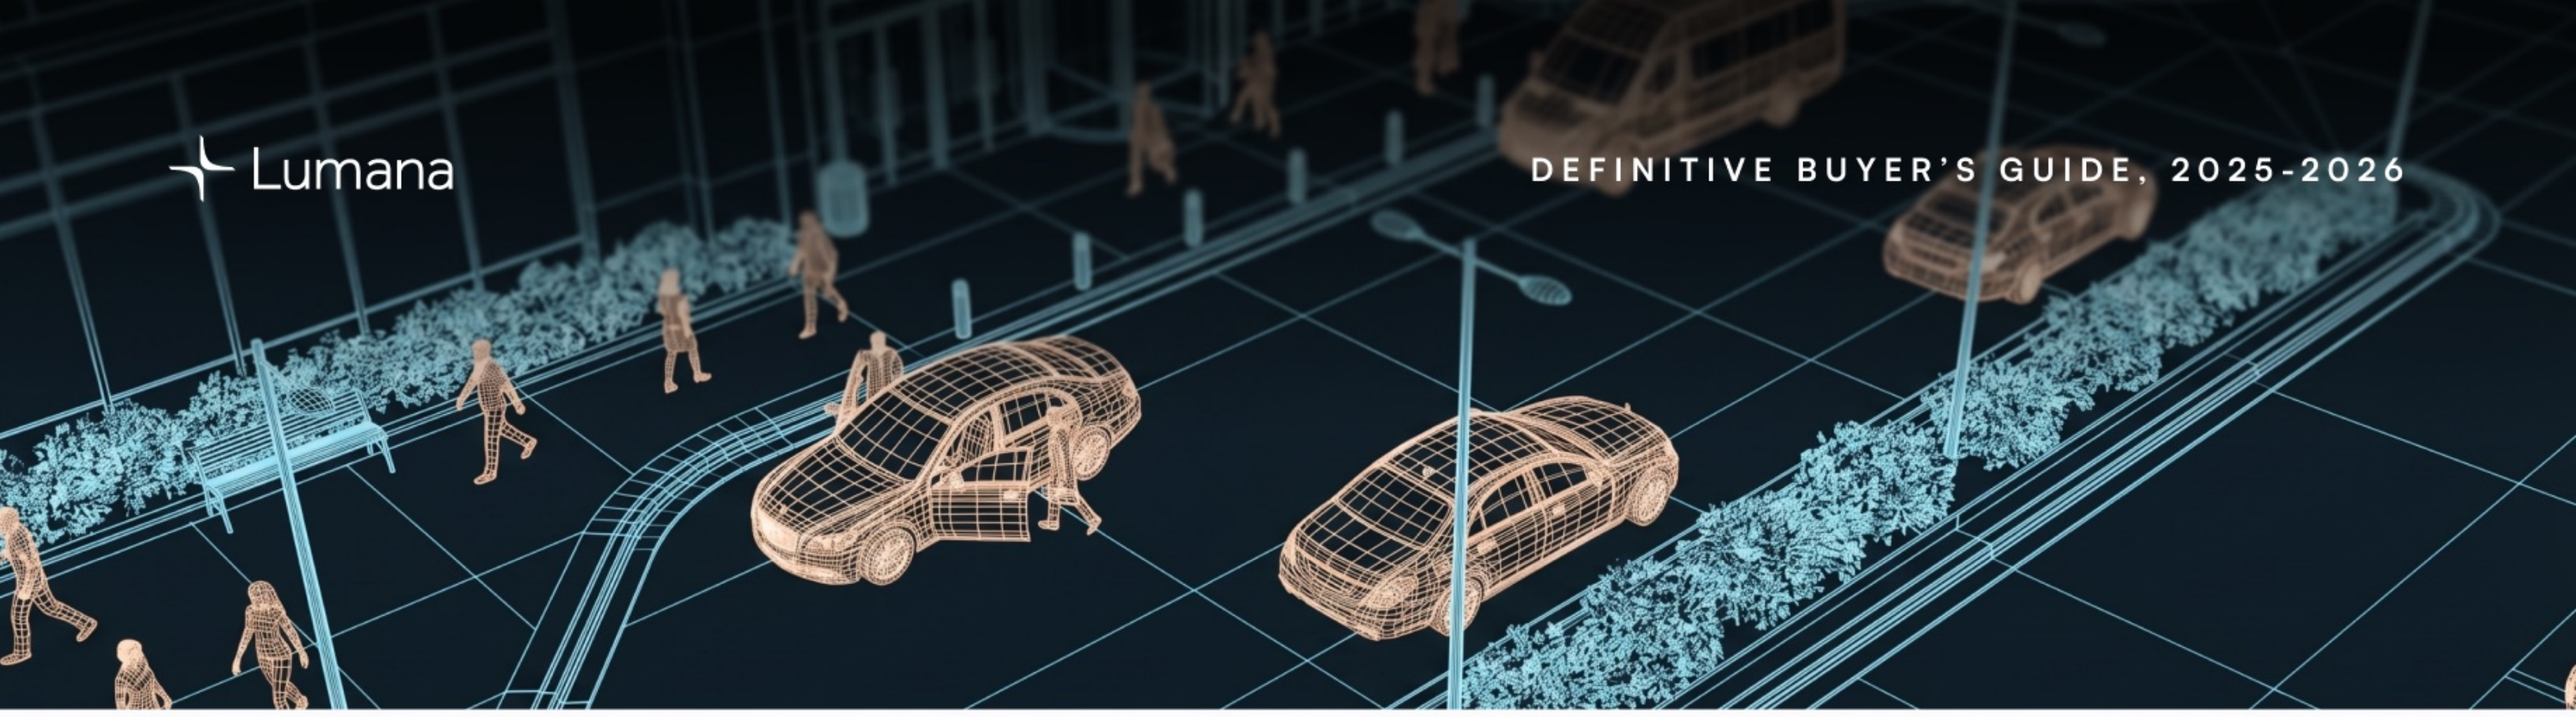

FACTOR 01

AI detection accuracy & continuous learning

Why it matters: Every security decision chains back to one question: can the system tell you what actually happened? Detection accuracy determines whether alerts are actionable or noise — and it separates platforms that market AI from platforms built on it.

The single most important differentiator in AI video security is not whether a platform uses AI — nearly every modern product makes that claim — but how the AI learns and whether it adapts to your specific environment over time.

Static AI models are trained on generic datasets and deployed without modification. They work reasonably well in controlled conditions but degrade quickly in real-world environments: shifting light, seasonal changes, new objects, or evolving activity patterns. The result is a flood of false positives that security teams learn to tune out — effectively making the system useless.

Continuous learning systems change this equation entirely. By incrementally training on local operational data, they improve accuracy over time at your specific sites without requiring costly retraining or vendor intervention.

The evolution across generations

Generation 1: Legacy

Motion-only detection

Triggers on any pixel change — leaves blowing, headlights sweeping across a wall, weather patterns. Security teams report spending hours chasing false alarms, leading to alarm fatigue and ignored genuine events. No ability to distinguish humans from vehicles from animals.

Generation 2: Cloud

Basic object classification

Can distinguish person vs. vehicle in controlled conditions. Pre-trained models work acceptably on install but don't adapt. Poor performance in edge cases: nighttime, weather, occlusion, unusual angles. Accuracy drops as environment shifts.

Generation 3: AI-Native

Context-aware continuous learning

Understands context — not just "there is a person" but what they're doing, where they should or shouldn't be, and whether it's anomalous. Models personalize to each deployment site through federated learning, improving detection and reducing false alarms by up to 90% over time.

What to ask vendors: Does your model update after deployment? How does it learn from my specific environment without requiring manual annotation? What is your false positive rate at 90-day vs. 30-day post-install?

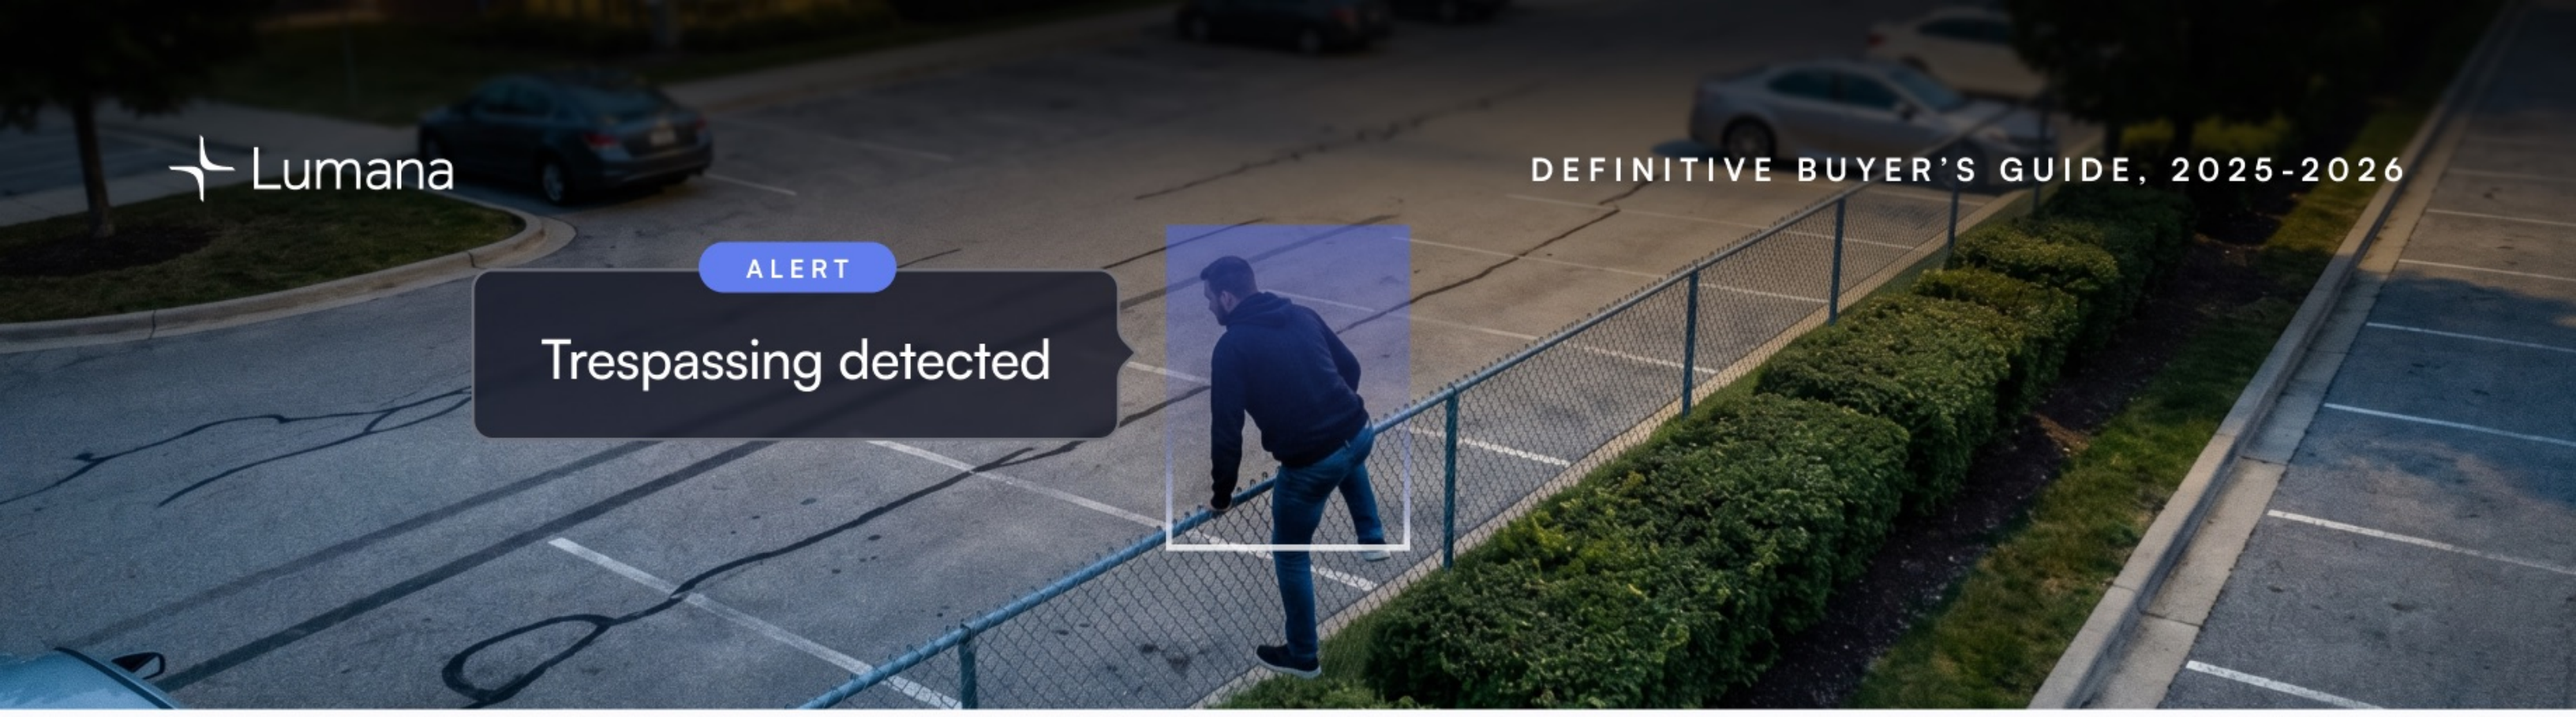

FACTOR 02

Real-time alerting & incident response

Why it matters: Security without timely response is documentation, not protection. The speed and specificity of alerts determines whether teams can act during an incident or are left reviewing what happened after it ends.

The gap between generation two and generation three video security is most visible in alerting. Cloud systems notify you that motion occurred. AI systems tell you that an unauthorized individual entered a restricted zone, stayed for 8 minutes, and then exited through the rear door — with the relevant footage surfaced automatically.

That specificity isn't cosmetic. It's the difference between a security team that responds to 200 alerts a day and one that responds to 12 meaningful ones. As organizations manage increasingly complex, multi-site environments, alert precision becomes an operational necessity.

> ***3.4M+*** *Detections processed at 99.99% object accuracy by Lumana during a single multi-venue event — enabling staff to act only on meaningful signals.*

The evolution across generations

Generation 1: Legacy

Reactive, after-the-fact

No automated alerting. Security events are discovered during manual footage review or physical patrol. Incidents are documented after they occur. Response time is measured in hours or days, not seconds.

Generation 2: Cloud

Motion notifications

Push notifications or emails when motion is detected. High noise-to-signal ratio. Teams receive hundreds of irrelevant alerts, leading to desensitization. Limited customization for zone-specific, time-based, or behavior-specific rules.

Generation 3: AI-Native

Intelligent, contextual alerts

Custom alerts built in natural language — "Alert me when a person enters the server room after 8pm" or "Notify if a vehicle remains in the loading dock for more than 15 minutes." Delivered via SMS, email, mobile app, Slack, or Teams with an attached video clip of the exact moment.

What to ask vendors: Can I write alert conditions in plain language without coding? What channels can alerts route to? What is the typical alert-to-review time from detection to notification?

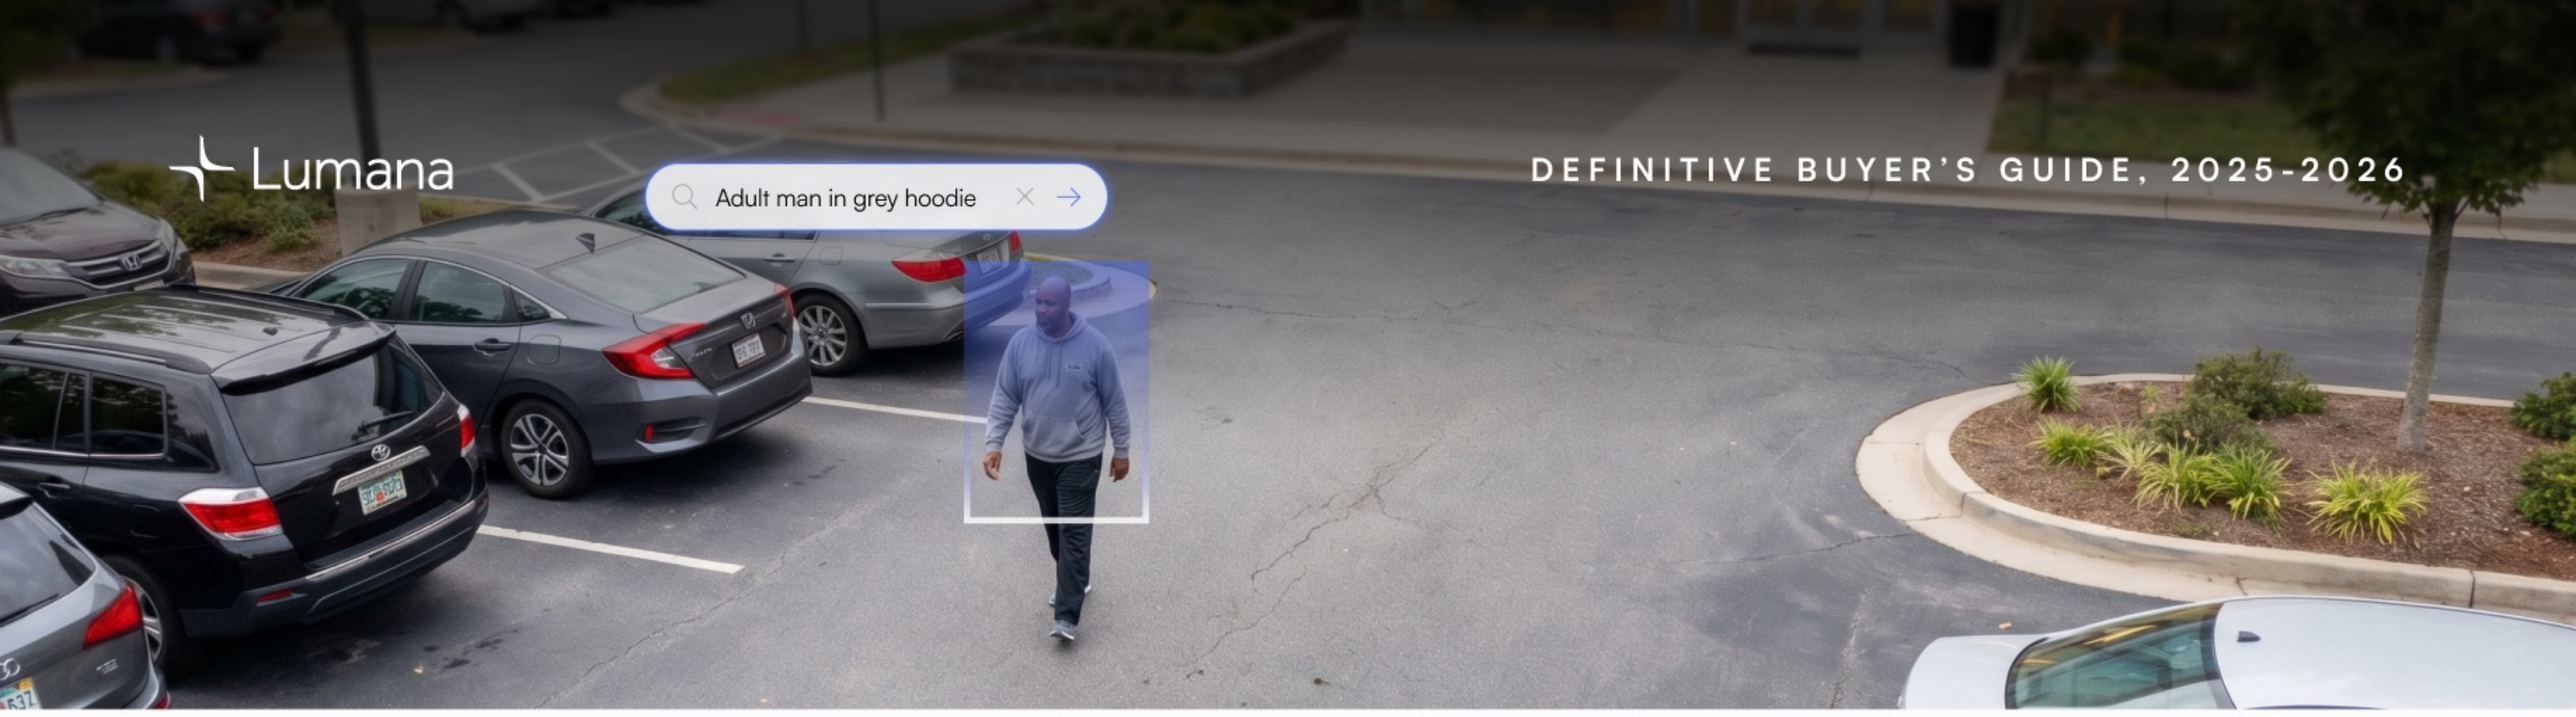

FACTOR 03

Intelligent search & investigation

Why it matters: Every security incident eventually requires investigation — whether for insurance, compliance, HR, or law enforcement. How quickly and accurately a system can surface relevant footage determines both response quality and operational overhead.

The hidden cost of traditional video security is investigator time. Reviewing hours of footage from multiple cameras to find a 30-second clip is standard practice in legacy environments. Studies suggest security teams spend 30—60% of their time on footage review — time that could be redirected to actual threat response.

AI search transforms this calculus. Instead of scrubbing timelines, investigators query the system the way they'd query a database: "Show me all instances of a person in a red jacket near the east entrance between 2pm and 4pm Wednesday." Results surface in seconds, not hours.

The evolution across generations

Generation 1: Legacy

Manual timeline scrubbing

Investigators manually review footage at 1x or 2x speed. Exporting evidence requires a physical USB drive. Multi-camera investigation means opening parallel windows and manually correlating timestamps. Hours become the baseline unit of investigation time.

Generation 2: Cloud

Timestamp & camera filtering

Remote access makes footage retrieval possible from anywhere. Filter by camera, date, and time. Some platforms offer basic motion-period filters. Still primarily a visual scrubbing workflow — just done remotely rather than on a DVR in the server room.

Generation 3: AI-Native

Natural language and attribute search

Search by object type, color, behavior, location, or time with natural language queries. AI automatically surfaces the most relevant clips across all cameras simultaneously. Cross-camera person tracking reconstructs a full subject journey across a property in seconds, not hours.

What to ask vendors: Can I search by appearance attributes (clothing color, vehicle type)? Can the system track an individual's movement across multiple cameras automatically? How does investigation time compare to your legacy platform benchmarks?

FACTOR 04

Scalability & camera compatibility

Why it matters: The true cost of any security platform is revealed when you grow. A system that requires forklift replacement of cameras every three years, or doubles in cost with each new site, is a liability — not an investment.

Organizations evaluating video security often overlook the long-term scalability penalty of closed-ecosystem platforms. Proprietary cameras lock you into vendor pricing, limit your ability to negotiate, and mean a complete replacement cost if you ever switch platforms.

Open, camera-agnostic platforms that work with existing IP cameras eliminate this friction. The analytics and AI ride on top of your existing infrastructure — protecting your prior capital investment while delivering the intelligence layer that legacy hardware alone could never provide.

The evolution across generations

Generation 1: Legacy

Hardware-bound, siloed

Each site requires dedicated NVR/DVR hardware. Adding sites means adding hardware, licensing, and IT complexity. Systems don't communicate across sites. Scaling to 50+ locations is an IT project measured in months, not days.

Generation 2: Cloud

Easier scaling, often proprietary

Cloud management simplifies multi-site deployment. However, many providers — including some market leaders — require proprietary cameras, creating a captive hardware ecosystem. Switching costs are high, and camera replacement is bundled into every upgrade cycle.

Generation 3: AI-Native

Camera-agnostic, unlimited scale

True AI platforms layer intelligence on top of any IP camera — your existing fleet, new hardware, different brands across different sites. Unlimited camera and user support means growth never triggers a platform change. New sites can be onboarded in days rather than months.

What to ask vendors: Does your platform require proprietary cameras, or is it camera-agnostic? What is the process for adding a new location? Are there per-camera licensing costs that compound at scale?

FACTOR 05

Cybersecurity & data privacy

Why it matters: Video security systems are themselves a security surface. Poorly secured surveillance infrastructure is an active liability — a vector for data breaches, unauthorized access, and regulatory non-compliance.

84% of security professionals cite cybersecurity concerns as a key factor in their technology decisions. That concern is well-founded. Legacy NVR systems are notorious for running unpatched firmware with hardcoded credentials. Cloud systems vary widely in their encryption and access control practices.

AI-native platforms built for the enterprise must clear a high bar: end-to-end encryption, SOC 2 compliance, multi-factor authentication, single sign-on, regular penetration testing, and automatic firmware updates. As biometric and AI privacy regulations expand globally — from GDPR to California's AI regulations — audit trails and data governance become non-optional.

The evolution across generations

Generation 1: Legacy

Significant exposure

Legacy NVR systems frequently run outdated firmware with known vulnerabilities. Hardcoded default credentials are common. No audit logs, no role-based access control, no encryption in transit. These systems are routinely exploited in botnet attacks.

Generation 2: Cloud

Baseline cloud security

Encrypted cloud storage and basic authentication. Some platforms offer 2FA and role-based access. Compliance posture varies significantly by vendor. Audit trails may be limited. Privacy controls often bolted on rather than built in.

Generation 3: AI-Native

Enterprise-grade, compliance-ready

SOC 2 Type II compliance, end-to-end encryption, MFA/SSO/SAML, role-based access, full audit trails, regular penetration testing, automatic security updates. Privacy-by-design architecture with configurable data retention, anonymization options, and regulatory compliance frameworks.

What to ask vendors: What is your SOC 2 compliance status? How is video encrypted in transit and at rest? Do you offer SAML/SSO integration? How are firmware updates managed, and what is your CVE response SLA?

FACTOR 06

Integration with existing systems

Why it matters: Security systems don't operate in isolation. Their value multiplies when they share context with access control, HR systems, sensor networks, and the communication tools teams already use every day.

Physical security has historically been fragmented by design: a video system here, an access control system there, alarm panels in a third interface. The overhead of jumping between systems during an active incident — while also managing a radio or phone — degrades response quality precisely when it needs to be highest.

Modern AI video platforms are built to unify these data streams. When a badge-in event triggers the video system to automatically surface footage from the relevant camera, security teams gain instant context without manually correlating logs. When a fight alert routes to Slack with a video clip attached, response can begin before anyone opens a separate security dashboard.

The evolution across generations

Generation 1: Legacy

Fully siloed

Video, access control, and alarms operate as independent systems. Integration requires custom development or third-party middleware. No native connection to communication tools. Events must be cross-referenced manually across systems during investigations.

Generation 2: Cloud

Limited integrations via APIs

Some open APIs for third-party access. Integration with access control systems is possible but may require configuration effort. Limited native connectivity to productivity tools. Webhook support varies. Often requires middleware or custom development for complex use cases.

Generation 3: AI-Native

Unified platform integrations

Native integrations with access control (Kisi, Genea), sensors (FLIR), communication (Slack, Microsoft Teams), and more. AI automatically correlates access events with video. Alerts route to the tools teams already use, with video context embedded. One platform, unified physical security intelligence.

What to ask vendors: What access control platforms do you natively integrate with? Can your system trigger responses in Slack or Teams? Do you support webhooks for custom integrations? Is your API documented and self-service?

FACTOR 07

Operational intelligence & analytics

Why it matters: The best ROI argument for AI video security isn't just fewer incidents — it's the operational intelligence that video generates continuously, turning a cost center into a source of business value.

This is where AI video security breaks entirely from the security-only paradigm. Once cameras can perceive and understand what's happening in a space, that intelligence becomes useful far beyond the security team.

Retail operations can understand foot traffic patterns and staff accordingly. Facilities teams can measure space utilization to optimize layouts. Manufacturing plants can monitor PPE compliance automatically. HR and EHS teams get visibility into workplace safety at scale without manual auditing. The cameras are already there — AI simply unlocks the value of the data they've always been generating.

The evolution across generations

Generation 1: Legacy

Zero operational value

Footage sits on a hard drive. The only people who access it are security investigators reviewing incidents. No aggregated analytics, no trend data, no business insights. The system is a pure cost with no operational upside.

Generation 2: Cloud

Basic counting & heatmaps

Some cloud platforms offer people counting, license plate recognition, and basic heat maps. These provide limited operational value but require separate licensing and lack the depth to drive meaningful business decisions. Analytics are often an add-on, not core.

Generation 3: AI-Native

Business intelligence layer

Real-time occupancy tracking, space utilization analytics, PPE compliance monitoring, vehicle flow analysis, queue management, anomaly detection for operational irregularities. Custom dashboards and scheduled reports translate video data into decisions that optimize staffing, space, safety, and operations.

What to ask vendors: Beyond security alerts, what operational analytics does your platform provide? Can I build custom dashboards without professional services? What use cases beyond security do your customers actively use the platform for?

FACTOR 08

Deployment, maintenance & total cost

Why it matters: The sticker price of a security system is rarely its real cost. Implementation timelines, IT overhead, ongoing maintenance, and hidden per-feature licensing determine whether an investment delivers returns or compounds expenses.

Legacy systems extract ongoing costs through hardware lifecycle replacement, on-site IT requirements, and maintenance contracts. Cloud platforms trade upfront hardware costs for recurring SaaS fees, but may layer per-camera, per-site, or per-feature costs that compound quickly at scale.

The total cost of ownership calculation for AI video security must account for the entire picture — and must also credit the operational savings on the other side: reduced false alarm response time, faster investigations, automated compliance monitoring, and security staff efficiency gains.

The evolution across generations

Generation 1: Legacy

High CapEx, high maintenance

Significant upfront hardware investment. On-site IT required for setup and ongoing maintenance. Hardware refresh cycles every 5—7 years. Physical service visits for troubleshooting. License and support contracts compound annually. No value outside of security.

Generation 2: Cloud

Reduced CapEx, OpEx scaling risk

Lower upfront investment, predictable subscription model. However, per-camera and per-feature pricing can grow significantly at scale. Proprietary hardware requirements can reset costs at upgrade cycles. Advanced analytics often require additional licensing tiers.

Generation 3: AI-Native

Works on existing hardware, compound ROI

Leverages existing camera infrastructure, eliminating hardware replacement costs. Fast setup (often within minutes per site). Automatic updates eliminate ongoing IT overhead. Operational analytics generate value that offsets licensing costs.

What to ask vendors: What is your per-camera cost at 100, 500, and 1,000 cameras? Are analytics features bundled or separate SKUs? What is the average time-to-value for a full deployment? What does your lifetime warranty cover?

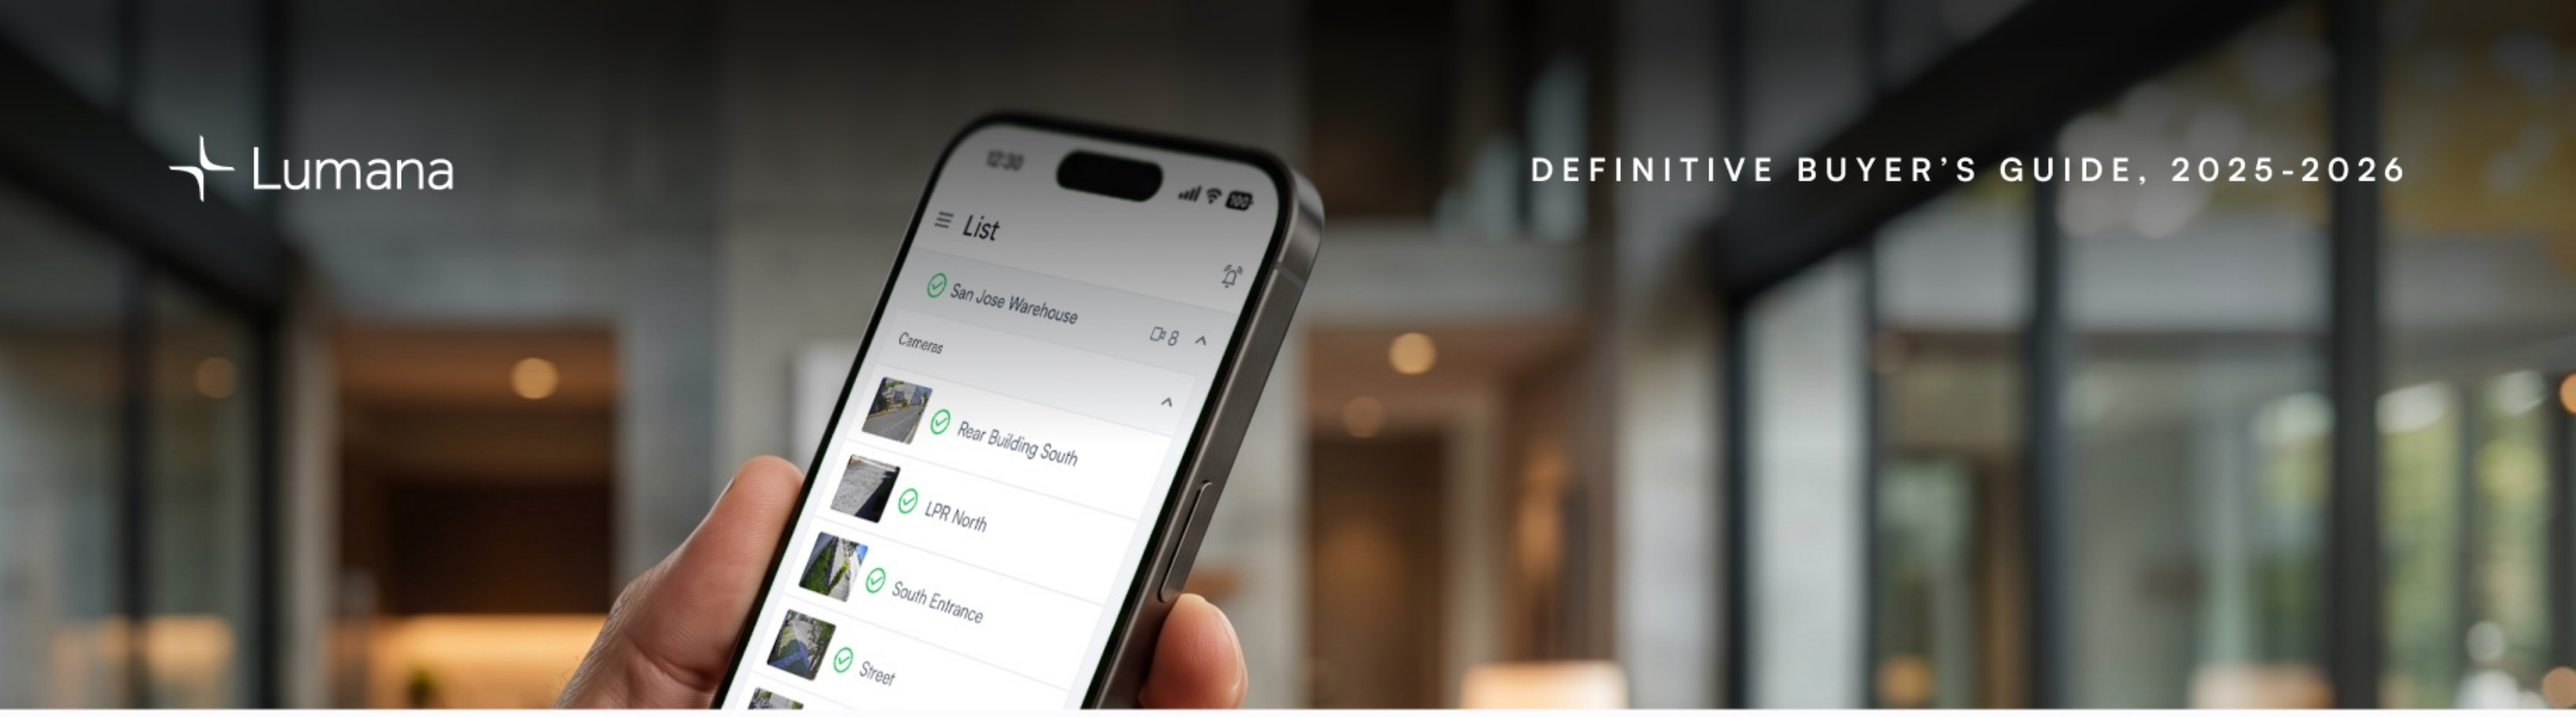

FACTOR 09

Reliability, uptime & support

Why it matters: Security infrastructure that fails during an incident isn't security infrastructure. Uptime guarantees, redundancy architecture, and support responsiveness are the operational backbone that everything else depends on.

Organizations with legacy systems know the experience of arriving at work to find a camera offline because the NVR filled up, or a door sensor failed to trigger a recording over the weekend. These gaps aren't inconveniences — they're liabilities.

Cloud-native and AI-native platforms move reliability management from on-site IT to vendor infrastructure with SLA-backed uptime guarantees. The best platforms go further: proactive monitoring that alerts your vendor to camera health issues before you notice them, and offline alerting that ensures no events are missed even during intermittent connectivity.

The evolution across generations

Generation 1: Legacy

Failure is common, silent

Hard drives fill up silently. Cameras go offline without alerts. The security team may not know footage was missed until an incident reveals the gap. On-site IT required for every troubleshooting event. No remote diagnostics.

Generation 2: Cloud

Remote monitoring, reactive

Cloud-managed systems provide remote visibility into camera status. Alerts when cameras go offline. Support teams can diagnose remotely. However, uptime monitoring may be reactive — you're notified of failures, not prevented from having them.

Generation 3: AI-Native

Proactive health monitoring, 24/7 support

Partner monitoring portals provide real-time, centralized diagnostics. Offline alerting via webhooks ensures events are never missed, even during connectivity gaps. 24/7 support with automatic software updates and no manual maintenance burden. Cloud backup with industrial SSD local storage provides redundancy.

What to ask vendors: What is your uptime SLA? Do you proactively alert me to camera health issues, or do I need to discover them? What happens to alerts if my connection is temporarily offline? What is your support response time commitment?

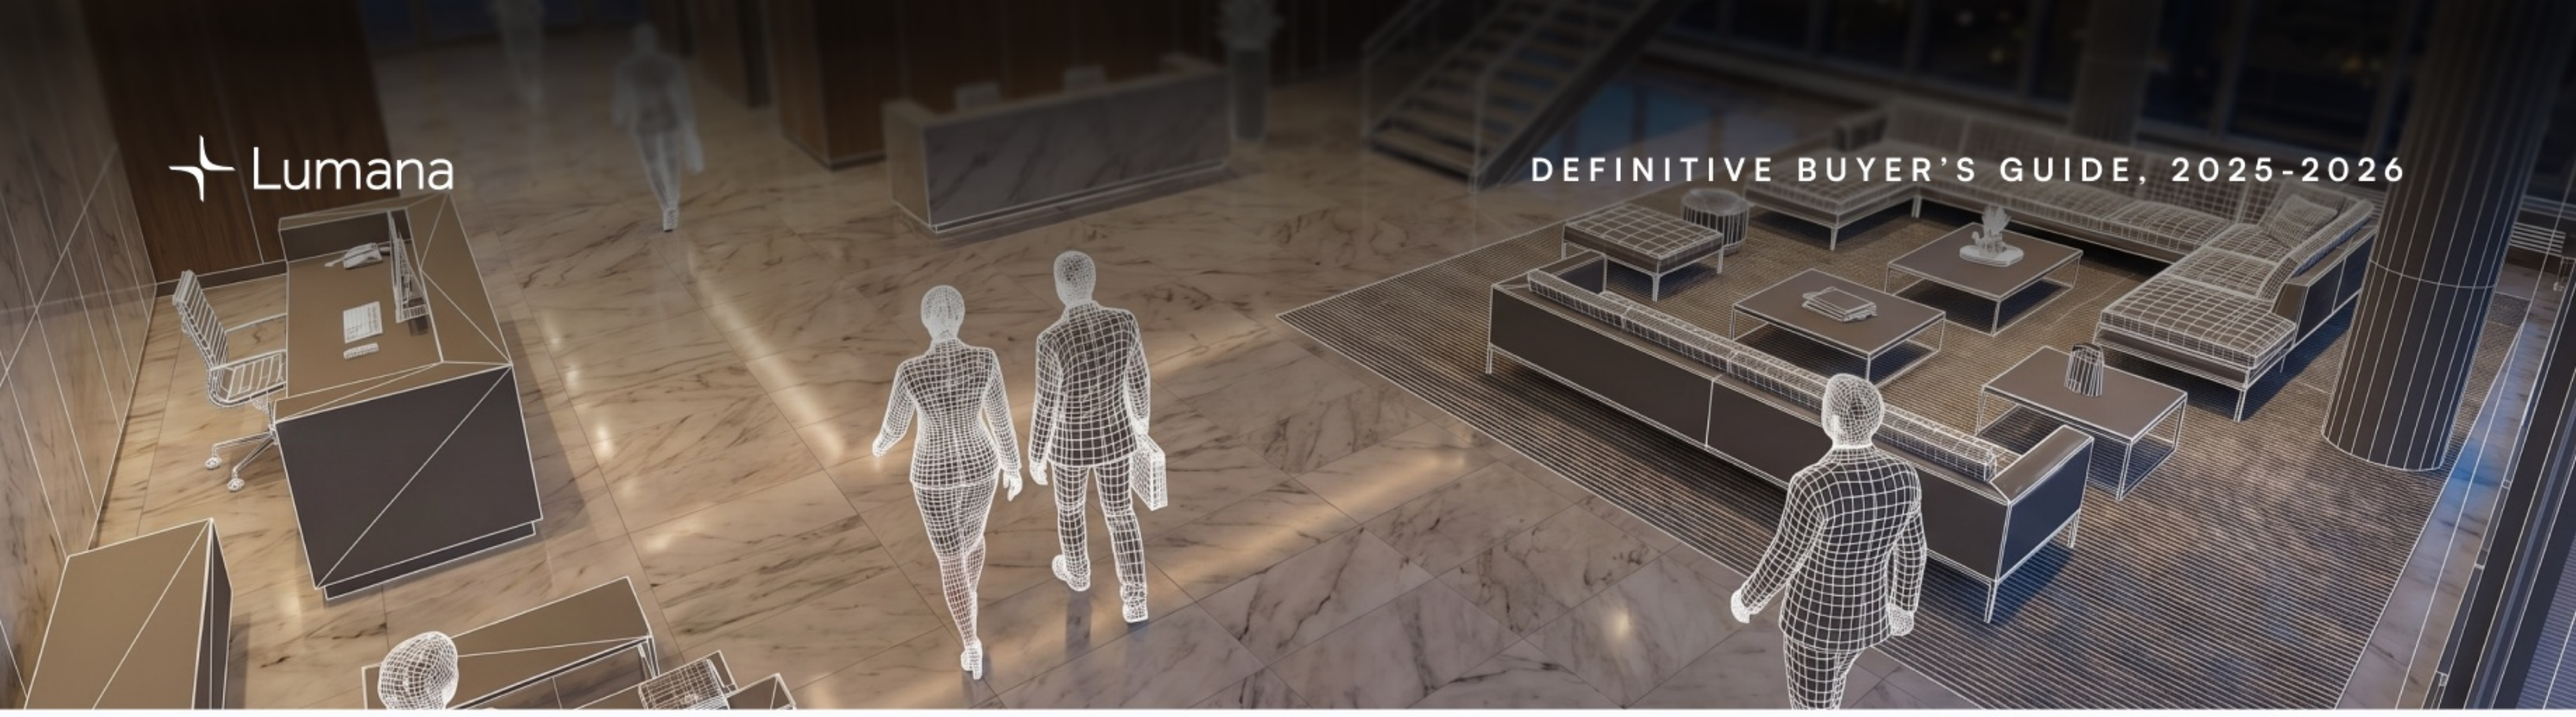

FACTOR 10

Vendor trajectory & roadmap

Why it matters: Security infrastructure is a multi-year commitment. The platform you deploy today will be responsible for protecting your organization in 2027 and beyond. Vendor momentum, funding, customer base, and product cadence indicate whether your investment will compound or depreciate.

The video security market is consolidating rapidly. Vendors that can't keep pace with AI advances will fall behind — and their customers will inherit the consequences. Legacy vendors are attempting to retrofit AI onto architectures that weren't built for it. True AI-native platforms continue to widen their capabilities advantage because intelligence compounds.

Evaluate vendors not just on where they are today, but on the evidence of sustained innovation: funding trajectory, product release cadence, customer adoption rates, and industry recognition. A vendor that has shipped 65 features in a single year and counts Meta, NYU, Salesforce, and McDonald's among its deployments is demonstrating a very different velocity than one maintaining a static feature set.

The evolution across generations

Generation 1: Legacy

Declining trajectory

Hardware-centric legacy vendors are losing ground to cloud and AI competitors. Product innovation is incremental. Customer base is largely locked in by switching costs, not value delivery. Investment in AI capabilities is retrofitted and limited.

Generation 2: Cloud

Consolidating

Cloud-native VMS providers are adding AI features to compete. Quality varies widely. Some platforms have strong AI investment and are genuine competitors; others are adding feature checkboxes. The category is differentiating, not commoditizing.

Generation 3: AI-Native

Accelerating innovation

Purpose-built AI platforms are expanding fastest: new model capabilities, deeper integrations, broader use cases, and expanding customer adoption. The AI video security market grows at 30.6% CAGR — platforms built for AI from the ground up are best positioned to capture and deliver that value.

What to ask vendors: How many features did you ship last year? What is your AI model update frequency? Who are your largest customers, and what use cases are they deploying? What is your public product roadmap for the next 12 months?

CHECKLIST

The AI Video Security Buyer's Checklist

Use these criteria in every vendor conversation and RFP process. A vendor confident in their platform will welcome all of these questions.

- [ ] **Continuous Learning AI** — Does the AI model adapt to my environment after deployment, or is it static?
- [ ] **False Positive Rate** — What is the documented false alarm rate, and how does it trend over time post-deployment?
- [ ] **Alert Customization** — Can I configure alerts in natural language without writing code?
- [ ] **Investigation Speed** — How long does it take to find specific footage using search, not timeline scrubbing?
- [ ] **Camera Compatibility** — Does the platform work with our existing IP cameras, or require proprietary hardware?
- [ ] **Compliance & Security** — SOC 2 certified? End-to-end encryption? MFA/SSO/SAML support?
- [ ] **System Integrations** — Native integrations with our access control system, sensors, and team communication tools?
- [ ] **Operational Analytics** — What business value does the platform provide beyond security incident detection?
- [ ] **Total Cost of Ownership** — Per-camera pricing at scale, plus hidden costs: analytics add-ons, support tiers, hardware requirements?
- [ ] **Uptime & Support SLA** — Documented uptime guarantee, offline alert capability, and 24/7 support response commitment?
- [ ] **Deployment Speed** — Realistic time from signed contract to first site fully operational?
- [ ] **Vendor Trajectory** — Feature release velocity, customer references at comparable scale, and public roadmap availability?

See what your cameras can actually tell you

Lumana works with your existing cameras to deliver AI video security that protects, informs, and improves your operations — from day one.

Request a demo at www.lumana.ai

50,000+
Cameras deployed

$64M
Total funding

#1
Best in Video Analytics, ISC West 2025

99.99%
Object Accuracy at Scale