



WHITE PAPER

The AI-Ready Building

An Informed Buyer's Guide to Future-Proofing Your Property's Technology

Why the security and communications platform you choose today determines which AI capabilities your building can — and can't — adopt tomorrow.

Audience	Condo boards, property managers, building owners, and the dealers who advise them
Reading time	Approximately 20 minutes
Published	September 2026

Contents

If this table appears empty, right-click it and choose "Update Field" (or press F9) — this happens automatically when the document is opened in Microsoft Word.

Contents..... 2

Executive Summary..... 3

Introduction: Why This Matters Now 4

Part 1 — AI in Plain English..... 5

Narrow AI (Machine Learning) 5

Edge AI 5

Cloud AI and Large Language Models (LLMs) 5

Agentic AI 5

Part 2 — The State of AI in Buildings Today 7

Access Control and Cameras 7

Resident and Tenant Communication 7

Portfolio Intelligence 7

Part 3 — The Hidden Cost of an AI-Blind Spec 9

Technical Debt: The Rip-and-Replace Trap 9

Trust Debt: AI Done Without Consent 9

Part 4 — AI Cuts Both Ways: The Cybersecurity Dimension 10

Part 5 — The AI-Readiness Framework 11

Part 6 — What This Looks Like Done Right 12

Conclusion..... 13

Appendix A — Glossary of AI Terms 14

Sources..... 16

Executive Summary

Every building technology purchase made today is, whether the buyer realizes it or not, an AI purchase. The access control platform, camera system, and communications backbone a board approves this year will determine which AI-powered capabilities that building can adopt three, five, or ten years from now — and at what cost. Buy the wrong foundation, and "adding AI later" quietly becomes "replacing the system later."

Two distinct technologies matter here, and they are frequently confused with each other. Large language models (LLMs) run in the cloud and handle language and reasoning: chatbots, report summarization, natural-language search across building data. Edge AI runs locally, on the camera or panel itself, and handles real-time perception: recognizing a face, reading a licence plate, flagging an anomaly, in milliseconds, without sending video off-site. A building — and the advisor recommending its systems — needs to understand both, because they solve different problems, carry different privacy implications, and depend on different hardware.

This paper explains both technologies in plain language, reviews the current state of AI adoption in buildings with sourced, current data, lays out the hidden costs of an AI-blind purchase, and provides a practical AI-Readiness Checklist that any board, property manager, or advisor can use to evaluate a proposal before signing it.

Introduction: Why This Matters Now

For most of the last decade, "smart building" meant connectivity: devices on a network, dashboards instead of clipboards, an app instead of a fob. That era is ending. The next decade of building technology is being defined by artificial intelligence — not as an add-on feature, but as the layer that decides how a building sees, responds, and communicates.

The scale of that shift is already visible in the numbers. The global AI market is on pace to exceed \$500 billion in 2026, growing at a compound annual rate of roughly 20.9%, and companies that have moved aggressively on AI adoption are seeing about 15% faster annual revenue growth than slower-moving competitors.

Smart buildings remain a small fraction of that spend today — an estimated 0.5% of global AI revenue in 2020, projected to reach only about 0.8% by 2028 — but more than 300 companies have already launched AI offerings specifically for building operators, and security and access control are the most mature category of deployment. The headroom is enormous, and it is closing faster than most procurement cycles move.

"At some point all software becomes AI because it will write and improve itself."

— James McHale, CEO, Memoori

That is the uncomfortable reality for anyone specifying building technology today: the system you choose is not a snapshot of what the building needs right now, it is a bet on what the building will be able to do for the next decade. A board that approves a closed, non-upgradable platform because it was the cheapest quote is not just buying cameras and door readers — it is quietly deciding that this building will not have facial recognition, predictive maintenance, or natural-language incident reporting unless it pays to rip the system out and start over.

That is why this paper treats AI literacy as a professional obligation, not a technical curiosity. An advisor — a dealer, an integrator, a consultant — who recommends a system without understanding its AI trajectory is not being neutral; they are making an AI decision on the customer's behalf, just an uninformed one. If you are not advising your customers on product fit for an AI-driven future, you are setting them up to fail — and, when the bill for a premature replacement eventually arrives, setting yourself up to fail with them.

Part 1 — AI in Plain English

"AI" gets used as a single word for at least four genuinely different technologies. Conflating them is where most confusing vendor proposals — and most confused buyers — come from. Here is the distinction that actually matters for a building.

Narrow AI (Machine Learning)

Narrow AI, sometimes just called machine learning, is software trained to do one bounded task well: tell a person apart from a vehicle, recognize a specific sound, flag a pattern that doesn't match the baseline. It has no general understanding of the world — it is very good at exactly one thing. Almost every "smart" feature in a modern security system, from false-alarm filtering to motion classification, is narrow AI under the hood.

Edge AI

Edge AI describes where the processing happens, not what kind of AI it is: the inference runs locally, on the device itself — a camera's onboard chip, a dedicated on-site appliance — instead of being sent to a data center. This matters enormously for real-time tasks like facial recognition or panic-alarm response, where a round trip to the cloud adds delay a building can't afford, and where sending continuous video off-site raises privacy and compliance questions a building may not want to take on.

Cloud AI and Large Language Models (LLMs)

LLMs are the technology behind tools like ChatGPT: models trained on enormous amounts of text that are good at language and reasoning — understanding a question, summarizing a report, answering "which units had a smoke alarm trigger this month" in plain English instead of a filtered spreadsheet. They typically run in the cloud, because they require far more computing power than a device sitting in a lobby closet can offer.

Agentic AI

The newest layer, agentic AI, doesn't just answer a question — it takes multi-step action on a person's behalf: reading a resident's maintenance email, creating the work order, assigning it to the right technician, and closing the loop, without a human retyping any of it. This is where LLM-style reasoning and real-world building systems start to merge.

Type of AI	What It Actually Does	Where It Shows Up
Narrow AI / ML	Pattern recognition on one bounded task	False-alarm filtering, motion classification
Edge AI	Real-time perception, processed on-device	Facial recognition, licence-plate reads, panic response
Cloud AI / LLMs	Language, reasoning, natural-language search	Resident chatbots, report summarization

Type of AI	What It Actually Does	Where It Shows Up
Agentic AI	Autonomous multi-step action across systems	Auto-generated work orders, alarm escalation


The practical takeaway: a building doesn't need to choose between these. It needs a platform that can run edge AI where speed and privacy demand it, and connect cleanly to cloud AI where language and reasoning add value — without requiring a hardware replacement to do either.

Part 2 — The State of AI in Buildings Today

AI in buildings is no longer a pilot project. It is already running in three areas that touch almost every resident interaction.

Access Control and Cameras

Security and access control are the most mature category of AI deployment in buildings today, encompassing facial recognition for entry, multi-factor biometric authentication, anomaly and behaviour detection, and automated visitor management. Facial recognition access specifically has already been adopted by more than one million users across major real estate companies. The barriers slowing wider adoption are practical, not technical: roughly a third of security professionals cite high implementation cost, and about a fifth cite unclear return on investment.



That latency gap is structural, not a marketing claim — it comes from eliminating the network round-trip, not from a faster chip. On-device hardware has caught up to make this practical almost anywhere: modern edge accelerators deliver serious compute in power-constrained packages, from roughly 26 TOPS (trillion operations per second) at 2.5 watts on entry-level modules, up to 275 TOPS at around 60 watts on high-end ones. The dominant real-world pattern is "edge inference with cloud orchestration" — detection and recognition happen locally, while standardized metadata flows to the cloud for cross-site dashboards and reporting, preserving both speed and privacy.

Resident and Tenant Communication

The clearest current use of LLM-style AI in buildings is communication. AI-driven leasing and communication assistants are already delivering measurable operational results: providers report that a large majority of operators using AI communication tools see moderate to significant reductions in operating expense, and a similar majority report measurable improvements in lead-to-lease conversion. On the maintenance side, natural-language processing is being used to turn an unstructured resident email directly into a structured, correctly categorized work order — with some platforms reporting close to a third faster resolution time as a result, and staff getting a third of their working time back from tasks AI now handles.

Portfolio Intelligence

A newer but fast-growing application is natural-language Q&A over a building's own operational data — a manager asking a plain-English question about leases, incidents, or maintenance history and getting a sourced answer back, instead of exporting a spreadsheet and searching it by hand. This is the clearest early example of agentic AI moving from concept to daily use in property operations.

Industry analysts describe the sector as still being in an early "narrow AI" phase — deep in specific use cases like access control and communication, not yet at general building intelligence. That phase will not last long. The gap between a building whose systems were built to add AI capability through software, and one that was not, is the entire subject of the next section.

Part 3 — The Hidden Cost of an AI-Blind Spec

There are two ways a building pays for an AI-blind technology decision. The first is financial: technical debt. The second is reputational: trust debt. Both are avoidable, and both are usually invisible at the moment of purchase.

Technical Debt: The Rip-and-Replace Trap

Security hardware and software typically carry a manufacturer support window of seven to ten years. Once that window closes, patches stop — and a system running an unsupported operating system cannot be secured against new threats no matter how well it was configured on day one. For a closed platform that also can't receive new AI capability through software, that same clock is effectively the building's AI adoption deadline: whatever the system could do when it launched is close to all it will ever do, and the eventual fix is a full, disruptive, expensive replacement rather than an update.

The alternative is not complicated, but it does need to be a deliberate purchasing criterion rather than an assumption: a platform whose hardware is designed to receive new capability — including AI-driven capability — through software updates, so the building's AI features improve on the same install rather than requiring a new one.

Trust Debt: AI Done Without Consent

Capability is only half of the equation. In 2019, tenants of the Atlantic Plaza Towers in Brooklyn successfully fought a landlord's proposal to install a facial recognition entry system, specifically over concerns about constant monitoring and how their biometric data would be stored and used. The system in that case may well have worked technically. It failed anyway, because privacy and consent were treated as an afterthought rather than a design requirement.

That is not a one-off cautionary tale; it is a preview of a live debate. Documented technical limitations — accuracy that degrades in low light, measurable bias affecting some demographics, and vulnerability to spoofing — mean that "the technology works" is necessary but not sufficient. A building considering biometric AI needs to be able to answer, before residents ask: where is this data stored, who can access it, and can it be deleted? A system that stores biometric templates locally, on-device, rather than in a centralized database most residents will never see, answers that question far more comfortably than one that doesn't.

Part 4 — AI Cuts Both Ways: The Cybersecurity Dimension

An AI-ready building is not only one that can adopt AI features. It is one that can defend against AI-enabled attacks — because the same advances making building AI more capable are making the threats against it more capable too.

More than 60% of organizations now name AI-based social engineering as one of their top security risks for 2026: phishing campaigns that convincingly impersonate IT staff to obtain administrator access, and automated bots that continuously scan the internet for exposed camera interfaces still running on default or weak credentials. Deepfake and synthetic media add a newer threat to physical security specifically — a fabricated video feed used to mask an intrusion, or AI-generated audio mimicking a trusted voice to talk a guard through a door. And video archives themselves have become a ransomware target: an attacker who compromises a recorder, server, or cloud account can encrypt months of footage and threaten to leak it, turning a security system into a liability in a single breach.

“Proactive security investment costs far less than remediating the damage and reputational harm of a cyberattack after the fact.”

— Security Industry Association

The practical response is not exotic — role-based access control, multi-factor authentication, encrypted video in transit and at rest, network segmentation that isolates cameras from other critical systems, and staff trained to treat video and audio as potentially fabricated rather than automatically trustworthy. What makes this an AI-readiness issue, and not just a general IT hygiene issue, is that it can't be solved once at installation. It requires an ongoing commitment — patches applied on a schedule, support that doesn't quietly lapse after year three — from whichever vendor is behind the platform. A purchase decision that doesn't include a stated, multi-year answer to "who keeps this secure, and for how long" is an incomplete purchase decision.

Part 5 — The AI-Readiness Framework

The questions below are the practical output of everything above. They are written to be asked directly of any vendor or dealer, on any building technology proposal, whether or not "AI" appears anywhere in the sales pitch.

A GOOD PROPOSAL WILL...	BE CAUTIOUS IF THE PROPOSAL...
Names whether AI inference runs at the edge, in the cloud, or both — and why.	Doesn't specify where AI processing happens, or treats it as a detail that doesn't matter.
Delivers new AI capabilities through software updates to existing hardware.	Requires a new hardware purchase for each new AI feature the building wants.
Includes ongoing cybersecurity patching and factory support under a stated, multi-year SLA.	Patching and support end quietly, with no stated lifecycle commitment.
Stores biometric templates on-device, with a clear, resident-facing data policy.	Centralizes biometric data with no stated retention, access, or deletion policy.
States a specific multi-year support lifecycle for both hardware and software.	Can't answer how many years the platform will remain supported.
Documents how core functions keep working if internet connectivity drops.	Depends entirely on cloud connectivity for core building functions.

None of these questions require a technical background to ask, and a vendor who can't answer them clearly and specifically is telling a board something important about the proposal, regardless of price. Boards and property managers are encouraged to bring this checklist to any technology proposal; dealers and advisors are encouraged to be able to answer every row of it before they present one.

Part 6 — What This Looks Like Done Right

Condoplex Systems was built around the premise this paper argues for: that a building's technology platform should be able to absorb new capability, including AI-driven capability, without being torn out and replaced. A few concrete examples of what that looks like in practice on the Condoplex platform:

- AI-assisted monitoring, facial recognition, and virtual guard services are on the Condoplex platform roadmap as capabilities delivered through software updates to compatible buildings — not as a future hardware upgrade cycle.
- The system is designed to stay fully manageable on site during an internet outage, including audio and video calls within the building, so core security functions never depend entirely on cloud connectivity.
- Condoplex is the only provider in its category that both provides and maintains end-to-end cybersecurity for the life of the product, with software updates and cybersecurity patches applied multiple times a year under a defined service agreement — not a one-time installation.
- Access credentials use encrypted Mifare Desfire EV3 technology rather than legacy unencrypted cards, and the platform's certified dealer network provides the ongoing local relationship this paper argues is a prerequisite for AI readiness, not an afterthought to it.

None of this makes the AI-readiness questions in Part 5 go away — if anything, it is the answer key. A board evaluating any proposal, Condoplex or otherwise, should expect a specific, confident answer to each row of that checklist. That standard, not any particular brand, is the actual point of this paper.

Conclusion

The building technology decision in front of a board today is bigger than it looks. It is not just a choice between camera brands or access control vendors — it is a decision about which version of the next decade's AI capability this building gets to have, and whether that capability arrives as a software update or as a disruptive, expensive replacement project.

The good news is that this decision doesn't require a technical background to get right. It requires asking the right questions, in the right order, before signing — the same six questions laid out in Part 5. Boards and property managers should ask them of every proposal. Dealers and advisors should be able to answer all six before they ever present one. Get that right, and AI readiness stops being a buzzword and becomes what it actually is: an ordinary part of a well-run building.

NEXT STEP

Bring the AI-Readiness Checklist in Part 5 to your next technology proposal, or speak with a certified Condoplex dealer about how these principles apply to your building specifically. condoplexinc.com · info@condoplexinc.com · (416) 674-9800

Appendix A — Glossary of AI Terms

Artificial Intelligence (AI)

Software that performs tasks which would normally require human perception, judgment, or reasoning — an umbrella term covering everything from simple pattern-matching to large language models.

Machine Learning (ML)

A method of building AI where a system learns patterns from data rather than following hand-written rules. Most practical AI in buildings today is machine learning.

Narrow AI

AI trained to perform one specific, bounded task well — such as telling a person apart from a vehicle — with no general understanding beyond that task.

Large Language Model (LLM)

An AI model trained on very large amounts of text, specialized in language and reasoning tasks: answering questions, summarizing documents, holding a conversation. The technology behind tools like ChatGPT.

Edge AI

AI processing that happens locally, on the device itself (a camera, panel, or on-site appliance) rather than being sent to a remote data center. Chosen for speed and for keeping sensitive data on-premises.

Cloud AI

AI processing that happens off-site, in a remote data center, typically because the task (like running a large language model) needs more computing power than a local device can provide.

Inference

The moment an already-trained AI model is actually used — for example, a camera analyzing a live frame of video to decide whether it shows a person. "Inference speed" is how quickly that decision happens.

TOPS (Trillion Operations Per Second)

A measure of how much raw computing power an AI chip can deliver — used to compare the capability of different edge AI hardware.

Agentic AI

AI that doesn't just answer a question but takes multi-step action on a person's behalf across one or more systems — for example, turning a resident's email directly into an assigned, tracked work order.

Biometric Template

A mathematical representation of a physical trait (a face, a fingerprint) used to verify identity. Distinct from a photo or fingerprint image — it is the encoded data a system actually compares against.

ONVIF

An industry standard that lets security devices from different manufacturers exchange data (such as event metadata) in a common format, rather than each vendor using an incompatible proprietary format.

Deepfake / Synthetic Media

Fabricated video, image, or audio generated by AI to convincingly impersonate a real person or scene — an emerging threat to the reliability of video surveillance as evidence.

CI/CD (Continuous Integration / Continuous Deployment)

A software practice of shipping small improvements and fixes on an ongoing basis rather than in large, infrequent releases. In a building context, it is what allows a security platform to gain new features — including AI features — through routine updates instead of hardware replacement.

Sources

This paper draws on the following third-party sources, current as of publication.

1. Edge vs. Cloud AI Inference in Physical Security: A Technology Briefing — Intellisee — <https://intellisee.com/intelligence/edge-vs-cloud-ai-inference-physical-security-technology-briefing/>
2. Tech Trends: How Edge AI Is Redefining the Camera's Role — Security Info Watch — <https://www.securityinfowatch.com/ai/article/55385289/tech-trends-how-edge-ai-is-redefining-the-cameras-role>
3. Biometric & Facial Recognition Access Control System Trends in 2026 — ButterflyMX — <https://butterflymx.com/blog/facial-recognition-access-control-system/>
4. 'Narrow' AI Adoption in Smart Commercial Buildings Expected to Grow — Security Info Watch — <https://www.securityinfowatch.com/integrators/article/53099754/narrow-ai-adoption-in-smart-commercial-buildings-expected-to-grow>
5. Legacy Systems: Rip and Replace or Keep Them Going? — Security Industry Association — <https://www.securityindustry.org/2022/10/14/legacy-systems-rip-and-replace-or-keep-them-going/>
6. Cybersecurity Threats to Watch Out For in 2026 — Network Optix — <https://www.networkoptix.com/blog/cybersecurity-threats-to-watch-out-for-in-2026>
7. Best AI-Powered Property Management Platforms: 2026 Guide — Re-Leased — <https://www.re-leased.com/software/best-ai-powered-property-management-platforms-2026-guide>

Smarter Systems. Safer Communities. Ready for What's Next.