Technology Community
A Community of IFMA

COGNITIVE CORP · STATE OF THE MARKET · JULY 2026

The Governed Building

The built environment has bought autonomy. It has not yet bought governance. That gap — not model performance — will separate the portfolios that lead the next decade from the ones that get audited.

78%
EXPECT AI TO RESHAPE REAL-ESTATE STRATEGY

15%
HAVE REACHED THE AI 'OPTIMIZATION' STAGE

$191B
INTELLIGENT BUILDING-AUTOMATION MARKET, 2030

$40M/7%
MAX AI ACT PENALTY (PROHIBITED PRACTICES)

COGNITIVE CORP · An evidence-benchmarked reading of the AI-in-buildings market, built on a public-source governance review

ABOUT THIS REPORT

A market reading of AI in the built environment — and the governance layer it is missing

This report examines a specific gap in the built-environment technology market: the distance between how fast AI is being deployed into building operations and how slowly the governance to control those decisions is being adopted. It draws its demand-side evidence from tier-one industry research — principally JLL's *Future of Work Survey 2026* (2,200+ leaders across 21 countries) and JLL's AI-in-real-estate research[1–3] — and its market-size and regulatory evidence from published analyst forecasts and the primary text of EU law[4–7]. Where the report makes claims about vendor governance, it does so from a public-source review and frames every absence as **"no public evidence that…"** rather than proof of non-existence.

The report is held to a higher evidentiary bar than the category it examines. Every quantitative and regulatory claim traces to a named source; the vendor review is reported on a public-source basis, with named platforms and the full evidence register open to qualified reviewers (Methodology). Its four criteria are grounded in external frameworks — NIST AI RMF, ISO/IEC 42001, and the EU AI Act — but their selection, scoring scale, and all-applicable-criteria threshold are Cognitive Corp's instrument: the concepts are external; the composite is ours.

Disclosure of interest. This report is published by **Cognitive Corp**, which sells AI-governance advisory for the built environment, including a proprietary framework (the "Building Constitution") and a Governance Gap Assessment. Cognitive Corp is therefore **not a neutral party** to its own thesis. Its framework appears in this report only as **one illustrative implementation** of governance principles that are themselves drawn from external standards; it is **self-attested**, was **not** assessed on the public-evidence basis applied to third-party platforms, and is **excluded from every comparative count** in this report. The author has no financial relationship with the third-party platforms referenced.

Contents

Key findings at a glance 3

Foreword — When pilots inherit authority 4

Methodology & evidence standard 5

01 The demand is real — and mispriced 7

02 The market bought autonomy, not governance 10

03 Governance stops being optional 13

04 From tools to capacity: the governed building 16

Implications · Limitations · References · Appendix 19

KEY FINDINGS AT A GLANCE

Five things the AI-in-buildings market has not priced in

78%
EXPECT AI TO RESHAPE RE STRATEGY (JLL 2026)[1]

15%
HAVE REACHED THE AI 'OPTIMIZATION' STAGE[1]

0 of 8
SELECTED PLATFORMS DOCUMENT EVERY APPLICABLE GOVERNANCE CRITERION†

Dec 2027
EU AI ACT HIGH-RISK OBLIGATIONS SCHEDULED TO APPLY[6]

01

Demand is broad; institutional maturity is not.

In JLL's 2026 survey of 2,200+ leaders, 78% expect AI to significantly reshape their real-estate strategy — yet only 15% have reached the "optimization" stage of redesigning roles and workplaces around it.[1] Intent is settled; the capacity to operate it is the gap.

02

Adoption is outrunning governance, not just maturity.

The intelligent-building-automation market — sensors, controls, services — is forecast to approach **$191 billion by 2030**, with AI increasingly embedded.[4] Yet in our public-source review of eight selected platforms, **none publicly documented every governance criterion applicable to its stated use cases**: strong security compliance, no published AI-governance layer.[†,8]

03

Security compliance is being mistaken for governance.

ISO 27001, SOC 2 and FedRAMP assure defined security controls; GDPR imposes data-protection duties and, in qualifying cases, safeguards around automated decisions. None alone establishes the four-part AI-governance capability assessed here — whether an agent's *decisions* are explainable, bounded, overseen, and tested. Buyers often buy the first and assume the second.[8,10]

04

Regulation converts the gap into liability.

The EU AI Act's high-risk obligations — risk management, data governance, logging, human oversight, transparency — are scheduled to apply from **2 December 2027** after the 2026 Digital Omnibus deferral. High-risk non-compliance may attract fines up to **€15M or 3%** of turnover (prohibited practices, up to **€35M or 7%**). Some building AI qualifies as high-risk where its intended purpose meets Article 6.[5,6,7]

05

The winning portfolios will build capacity, not buy tools.

Point tools automate a task; they do not make an organization able to govern it. The differentiator is an operating capability — explainability, human-in-the-loop, bias auditing, pre-deployment testing — anchored to external standards and owned by the operator.[1,2,9]

† "0 of 8" = our public-source review of eight selected platforms (July 2026), scored on the criteria on p.5; none publicly documented every applicable criterion. "No public evidence" is not proof of non-existence; named platforms and the register are open to qualified reviewers (Disclosure, p.2).

FOREWORD

When pilots inherit authority

For a decade, artificial intelligence in the built environment lived in the pilot. A pilot is a comfortable place. It has a fixed scope, a sympathetic sponsor, and no real authority over the building. If it misbehaves, someone turns it off and writes a lessons-learned memo. The stakes are a slide, not a lawsuit.

That is changing — not because pilots have ended, but because they are being handed authority. In JLL's 2026 survey, 78% of leaders expect AI to significantly reshape their real-estate strategy, even as only 15% have reached the stage of actively redesigning roles and workplaces around it.[1] The distance between that intent and that institutional maturity is exactly where the new risk lives. When AI moves from the pilot to the plant room, it stops recommending and starts deciding — dispatching work orders, throttling HVAC, allocating space, controlling access. The question changes with it. It is no longer "is the model any good?" It is "who governs what the model decides, and can they prove it?"

This report argues that the market has answered the first question and skipped the second. Vendors compete on autonomy — more automation, more agents, more decisions taken without a human in the loop. Few publicly demonstrate governance: the ability to explain a decision after the fact, to bound it before it is taken, to keep a human in the loop where the stakes are high, and to test the agent for safe behavior before it is given write access to a building. That is not a subtle omission. It is the difference between a system you can defend to a regulator, an insurer, a tenant, or a court — and one you cannot.

Our posture throughout is evidentiary. Demand and market figures come from named, tier-one sources. Regulatory claims are drawn from the primary text of EU law and its 2026 amendments, verified current as of publication. Where we assess the vendor field, we publish the yardstick, map it to external standards, and record every absence as "no public evidence" rather than a verdict. And because Cognitive Corp is not a disinterested observer of this thesis, our own framework is disclosed, walled off, and excluded from every count. The argument should stand on the industry's data, not on ours.

METHODOLOGY & EVIDENCE STANDARD

How this was built — and how to check it

Scope. The report covers AI used in the operation of commercial buildings and portfolios — facility management, workplace, and corporate real estate. Demand evidence is drawn principally from JLL's *Future of Work Survey 2026* (2,200+ C-suite and CRE leaders across 21 countries, fielded January–April 2026), with earlier editions and JLL's AI-in-real-estate research for trend and context.[1–3] Market sizing is from a published analyst forecast, cited by edition and date.[4] Regulatory analysis is confined to the EU AI Act as the most consequential in-force regime for building AI.[5–7] A platform qualified for the vendor review if it publicly markets AI that *takes* operational building decisions — or can be configured to act on them without instance-level human approval (write access). The reviewed set spans decision-support and autonomous-action platforms; findings are reported against each platform's *stated* autonomy (see definitions, right).

Instrument. Vendor governance is assessed on four criteria, each mapped to an external framework rather than a Cognitive Corp definition: *explainability* of autonomous decisions (NIST AI RMF Explainability; EU AI Act Art. 13), *human oversight* and enforceable decision boundaries (EU AI Act Art. 14; ISO/IEC 42001 Annex A.9), *bias / fairness assessment* of operational decisions (NIST AI RMF Fairness; ISO/IEC 42001 cl. 8.4), and *pre-deployment testing* of agent behavior (NIST AI RMF Measure/TEVV; ISO/IEC 42001 Annex A.6.2.4, verification & validation). The full crosswalk — and a caveat on clause-level precision — is in the Appendix. The individual criteria are externally grounded; their selection, operationalization, the scoring scale, and the all-applicable-criteria threshold are Cognitive Corp's instrument, not a composite published elsewhere.

Scoring. Each criterion is scored from public sources only, on a fixed scale:

Score	Definition
Present	Platform publicly documents the item as a governance capability (not merely a security or privacy control)
Partial	Item is referenced or adjacent (e.g., dashboards, audit logs) but not framed or operated as governance
Absent	No public evidence of the item as of the capture date — *not* proof of non-existence
Not applicable	The platform's documented intended purpose does not involve decisions capable of the harm the criterion addresses (e.g., person-level fairness for a setpoint-only optimizer); rationale recorded
Unknown	Insufficient public information to score

The evidence standard — and how qualified reviewers can audit it

The vendor review drew on ~40 public sources across the eight platforms' own documentation, product pages, trust/compliance centers, and analyst coverage, captured June–July 2026. Each scored cell records the platform, criterion, score, source URL, and capture date; a sample of fully-evidenced cells — including honest "Absent" rows — is

published in the Appendix. Any single claim can be challenged on its specific cell; the source is on record. Demand and regulatory figures are independently verifiable against the cited primary sources.

Definitions used in this report

AI system — software that, for a given objective, infers from inputs how to generate outputs (predictions, recommendations, decisions) per the EU AI Act's definition; a fixed deterministic control sequence is *automation* but not necessarily an AI system. **Autonomous decision** — an output acted upon without a human approving that specific instance. **Agent** — an AI system that can take actions in a building system toward a goal, not merely predict. **Write access** — authority to change a building system's state (setpoints, dispatch, access), as distinct from read-only analytics. **Consequential decision** — one with material effect on safety, cost, comfort, access, or equity. **Governance framework** — a published, operative set of controls for explainability, human oversight, bias assessment, and pre-deployment testing of such decisions — distinct from security, privacy, or quality certifications.

THE REVIEW PIPELINE — SELECTION TO REVIEW

Each cell records platform, criterion, score, source URL and capture date; criteria not applicable to a platform's purpose are marked Not applicable.

Counts of vendor governance exclude Cognitive Corp's own framework (self-attested; see Disclosure, p.2). Market-size and macro figures reported by third parties are attributed, not restated as independent fact. Governance positions are directional readings of public evidence, not certifications, and nothing here is legal advice.

01

FINDING ONE

The demand is real — and mispriced

Organizations have decided to deploy AI into building operations. What they have not yet decided is who controls the decisions it will make.

01 · THE DEMAND IS REAL

Adoption is no longer the question

The most important thing JLL's research establishes is that the built-environment's AI debate has moved past appetite. In the 2026 *Future of Work Survey* of 2,200+ leaders, 78% expect AI to significantly reshape their real-estate strategy, and 89% of C-suite leaders in JLL's AI research say AI can help solve major commercial real-estate challenges.[1,2] This is not a market deciding whether to move. It is a market that has decided to move — and is now discovering how much harder it is to operate AI than to want it.

The workforce backdrop reinforces the direction rather than undercutting it. In the 2026 survey, 60% of leaders expect their workforce to grow and a majority expect AI to reinvent human roles rather than replace them — an expansion story, not a contraction one.[1] Office-presence policies are tightening in parallel, reversing the fully-remote drift of the early 2020s.[1] More people, more space, and more AI in the loop — with a mandate to run it well, and not yet the capacity to govern it.

EXHIBIT 1 STRATEGIC EXPECTATION IS WIDESPREAD; INSTITUTIONAL MATURITY IS NOT

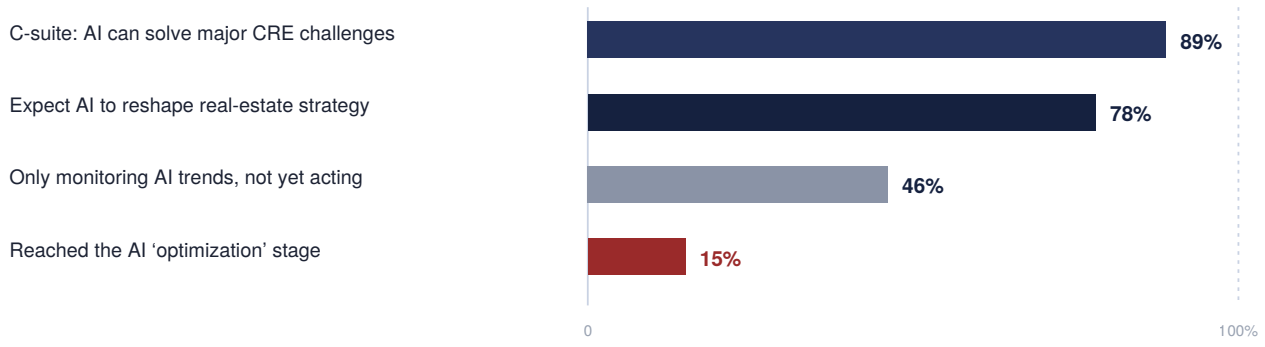

Share of leaders / organizations at each position. Sources: JLL Future of Work Survey 2026 (reshape strategy; maturity stages) and JLL AI-in-real-estate research (C-suite belief)[1,2].

What makes the demand *mispriced* is the gap the exhibit makes visible. Intent is widespread; operational maturity is scarce — only 15% of organizations have reached the "optimization" stage of actively redesigning around AI, while nearly half are still merely monitoring.[1] Where AI *is* being pushed into production, it is often being promoted without the operating scaffolding that production requires. The budget is being sized for capability — agents, automation, throughput — and not for control. That is the mispricing this report is about: the market is paying for what AI can *do* and assuming, for free, the ability to govern what it does.

> *Adoption risk has quietly inverted. The danger is no longer moving too slowly on AI — it is scaling autonomy faster than the capacity to govern it.*

01 · THE DEMAND IS REAL

The clock speed of the building is changing

Demand is not only growing; it is accelerating the metabolism of the organizations that hold real estate. JLL's interviews with senior executives at AI-native firms describe business cycles compressing from the traditional three-to-five-year planning horizon to **12–24 months**.[3] A portfolio strategy built for a five-year cadence now has to answer to a two-year one. Long leases, fixed fit-outs, and hierarchical floorplates were designed for a slower clock; they increasingly misalign with how AI-driven organizations actually work.[3]

This matters for governance because speed is the enemy of unexamined autonomy. When the building's systems are re-tuning continuously — reallocating space, re-sequencing maintenance, adjusting environmental setpoints — the volume of consequential decisions taken without direct human involvement rises sharply. Each one is small. In aggregate they determine who gets which desk, which tenant complaint is actioned first, which equipment is run to failure, and how much energy a floor consumes. At a two-year clock speed, an ungoverned decision engine can drift a long way before anyone notices.

Three roles the AI-native building creates

JLL's executives describe the workforce reorganizing around AI into *system thinkers* (who set the strategy the agents pursue), *system builders* (who design the human–AI workflows), and *system operators* (who monitor agents with human oversight).[3] Every one of those roles is a governance role in disguise — and none of them can function without a record of what the AI decided and why.

The strategic implication is uncomfortable for a market selling autonomy as the product. The organizations best positioned for a faster clock are not the ones with the most autonomous AI. They are the ones that can let AI move quickly *because* they can see, bound, and correct what it does. Governance is not the brake on this transition. It is the steering.

02

FINDING TWO

The market bought autonomy, not governance

Many leading platforms now ship autonomous or near-autonomous decision-making. In our public-source review of eight selected platforms, none publicly documented every governance criterion applicable to what those agents decide.

02 · AUTONOMY WITHOUT GOVERNANCE

A $191 billion market, and a governance gap

The money is moving, but it is worth being precise about what it measures. The broader intelligent-building-automation market — sensors, controls, connectivity, computing and services — is forecast to approach **$191 billion by 2030** (from roughly $122 billion in 2026; about 11.8% CAGR).[4] That figure is not a measure of autonomous-decision AI specifically; it is the envelope inside which such AI is increasingly embedded. Three things are rising together and should not be conflated: the building-automation market's value, the adoption of AI-enabled operations, and the deployment of systems with genuine autonomous decision authority. This report is about the third — and the governance layer that should accompany it.

What the market has not publicly demonstrated at the same pace is that governance layer. We reviewed eight smart-building AI platforms — spanning workplace analytics, autonomous HVAC optimization, connected-portfolio and work-order platforms, and building operating systems — against four governance criteria mapped to external standards (p.5). The pattern was consistent: strong security and privacy compliance, and **no public evidence**, as of July 2026, that any of the eight publicly documents every governance criterion applicable to its stated autonomous use cases.[8]

EXHIBIT 2 THE GOVERNANCE GAP: AUTONOMY (X) VS. PUBLISHED GOVERNANCE MATURITY (Y)

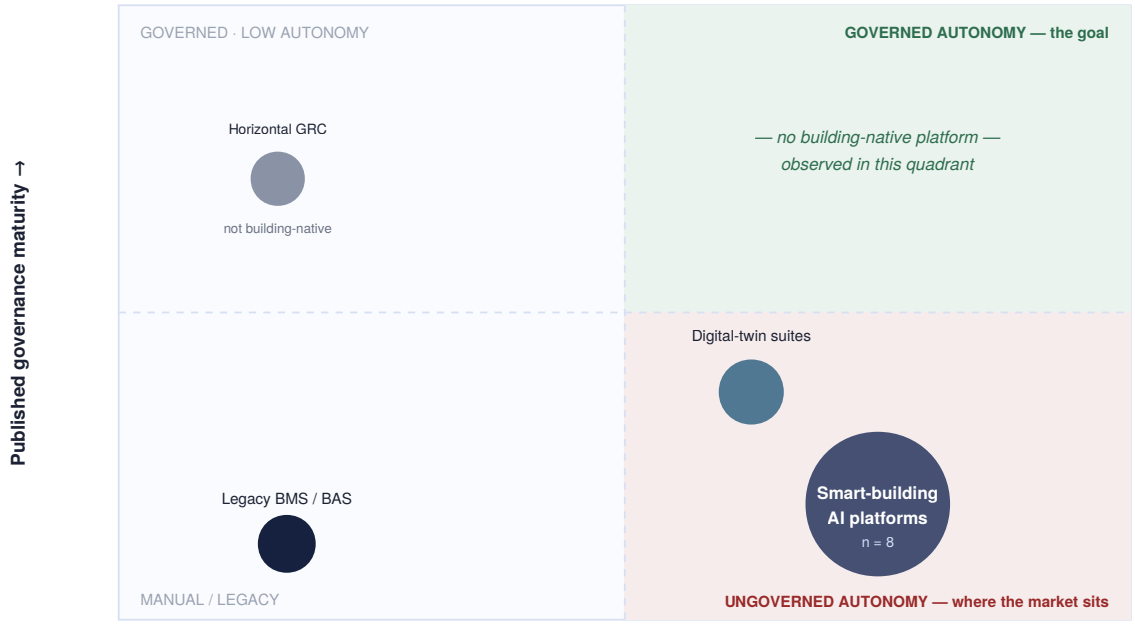

Bubble size = number of platforms in group; position = directional reading of public evidence (not an item-level ranking).

The empty top-right quadrant is the finding. Cognitive Corp's own framework is disclosed on p.2, excluded from all counts, and not plotted here.

Directional group positions from a public-source review of eight selected platforms, captured June–July 2026; named platforms available to qualified reviewers (p.5). Source: Cognitive Corp governance review[8].

The exhibit tells the whole story in one frame. Capability runs to the right; the market clusters in the bottom-right — high autonomy, low published governance. The top-right quadrant, where autonomy is matched by the ability to govern it, is effectively empty of building-native platforms. That emptiness is the finding. A published governance-framework layer — such as the one Cognitive Corp sells, disclosed on p.2 — is the category that would occupy that space; it is excluded from every count here, and named only to identify the missing layer, not to rank it.

02 · AUTONOMY WITHOUT GOVERNANCE

Security compliance is not governance — and the market is conflating them

The most important distinction in this report is the one buyers most often miss. When a vendor answers a governance question with a security certification, the answer sounds responsive but addresses a different problem. ISO 27001, SOC 2, and FedRAMP provide assurance over specified security controls within a defined scope; GDPR is a data-protection *law*, not a certification, and in qualifying cases it can itself require transparency about automated decisions, meaningful human intervention, and a right to contest them.[10] What none of them establishes on its own is a complete AI-governance system for what an autonomous agent *decides* — whether its decisions can be explained, bounded, overseen, or tested before the agent is trusted with a building.[8]

Security, privacy & assurance controls answer…	AI governance answers…
Is the data encrypted and access-controlled?	Can the agent **explain** why it made a specific decision, with the data and alternatives it weighed?
Can we resist and detect intrusion?	Is a **human in the loop** for irreversible or high-stakes actions, by enforceable rule?
Are we compliant with privacy law?	Are space, dispatch, and resource decisions **audited for bias** against tenants and staff?
Is the vendor audited (SOC 2 / ISO 27001)?	Was the agent **tested for safe behavior under pressure** before it got write access?

The columns are not substitutes: an organization can hold every certificate on the left and still answer none of the questions on the right — precisely what a regulator, an insurer, or a plaintiff's lawyer probes first. Certification of the pipe says nothing about the judgment of the water flowing through it.

> *The buildings that lead the next decade will not be the ones with the most autonomous AI. They will be the ones where someone thought to ask: who governs what the AI decides — and can we prove it?*

None of this indicts the platforms; they are excellent at what the market rewarded. The gap is a market failure, not a vendor one — and the next pressure comes not from buyers but from law.

03

FINDING THREE

Governance stops being optional

Regulation converts a competitive gap into a compliance liability — with a deadline and a penalty attached.

03 · GOVERNANCE BECOMES LAW

The EU AI Act can reach into the plant room

Until recently, the case for AI governance in buildings was a prudential one — wise, but discretionary. The EU AI Act removes the discretion for a meaningful class of building systems. The Act, in force since 1 August 2024, regulates AI by risk tier, and its **high-risk** obligations are the ones that reach building operations: risk-management systems, data governance, technical documentation, automatic logging, human oversight, transparency, and accuracy and robustness testing before a system is placed on the market.[5] Read that list against Exhibit 2's empty quadrant, and the problem states itself: the Act mandates precisely the capabilities the market has not built.

Which building systems are in scope? High-risk status is not automatic; it turns on a system's *intended purpose* and whether it meets the Act's Article 6 criteria via Annex I or Annex III. Some building AI can qualify — particularly where it performs a safety-critical function, such as AI governing life-safety, emergency access, or environmental controls in regulated settings. Ordinary comfort optimization, energy efficiency, predictive maintenance, or workflow automation is **not** high-risk merely because it runs in a building; the 2026 Omnibus text further clarifies that systems used only for performance optimization, efficiency, convenience, or quality control are not necessarily safety components unless their failure could endanger health or safety.[5,6] Separately, some uses are prohibited outright — including emotion-recognition systems in the workplace, a live question for occupancy and "wellness" analytics that infer employee affect.[5] The penalties are not nominal: up to **€35M or 7% of global turnover** for prohibited practices, and **€15M or 3%** for high-risk non-compliance.[7]

EXHIBIT 3 EU AI ACT: THE COMPLIANCE HORIZON FOR BUILDING AI

*High-risk (Annex III) obligations were deferred from 2 Aug 2026 to **2 Dec 2027** by the 2026 Digital Omnibus (Council final approval 29 June 2026); Annex I product-embedded AI follows 2 Aug 2028. Dates are subject to publication in the Official Journal. Sources: Regulation (EU) 2024/1689; Council of the European Union[5,6].*

A deferral is not a reprieve. In May 2026 the EU agreed a Digital Omnibus package that moved the high-risk deadline from August 2026 to **2 December 2027**, with the Council giving final approval on 29 June 2026.[6] The extra sixteen months are for building the compliance infrastructure, not for ignoring it — and for a portfolio that must instrument its systems, generate decision records, and stand up human-oversight workflows across dozens of buildings, sixteen months is a normal implementation runway, not a long one.

03 · GOVERNANCE BECOMES LAW

Why the deadline is closer than it looks

Regulation rarely arrives alone, and the AI Act is the leading edge, not the whole front. It applies extraterritorially to anyone placing a covered system on the EU market, so for a multinational occupier or global property manager, “we don’t operate in Europe” is rarely a complete answer — and a shrinking one.

The Act does not place the same duties on everyone in the chain, and the distinction matters for who does what. **Providers** — those who develop, brand, or substantially modify a high-risk system — carry the heaviest obligations. **Deployers** — those who use the system under their own authority, which is where most building owners and FM operators sit — carry a different and lighter set.[5,10] The catch for operators is Article 25: substantially modifying a system, or putting it out under your own name, can turn a deployer into a provider and pull the heavier duties across with it.

Provider responsibilities	Deployer / operator responsibilities
Conformity assessment, quality-management system, technical documentation, registration	Use per the provider’s instructions; assign competent human oversight
Lifecycle and post-market monitoring obligations	Monitor operation; keep the required logs; ensure input data is relevant
Applies when developing, branding, or *substantially modifying* the system (Art. 25)	Incident escalation and cooperation; a fundamental-rights impact assessment in defined cases

Most operators sit on the right-hand column, which is lighter than the compliance panic often implies. The catch is that a deployer’s duties — oversight, monitoring, logging, human-in-the-loop — are exactly the governance capabilities this report argues for. An operator that has closed the governance gap has already built them; one that bought only autonomy starts from lower apparent readiness — on a clock.[5,10]

The insurance and litigation flank

Regulators are not the only informed, hostile readers of an ungoverned decision. Insurers may increasingly examine AI-driven operational risk as loss experience and underwriting practices develop, and the first question after an incident — a wrongful access denial, a maintenance decision that preceded a failure, a biased space allocation — will be whether the operator can produce a decision record showing what the system did and why. “The vendor is SOC 2 certified” is not an answer to that question. A governed building can produce the record. An ungoverned one produces a shrug.

> *December 2027 is not a snooze button. For a portfolio, it is roughly one budget cycle to instrument, document, and staff the oversight the law will assume you already have.*

04

FINDING FOUR

From tools to capacity: the governed building

Governance is not a product you buy once. It is an operating capability — four requirements, each anchored to an external standard, owned by the operator.

04 · FROM TOOLS TO CAPACITY

What a governed building actually requires

If the gap is real and the deadline is fixed, the constructive question follows: what does it take to close it? The answer is not another autonomous tool. Tools automate tasks; they do not confer the organizational capacity to govern what has been automated. Closing the gap means standing up four capabilities and operating them continuously — each one addressable, testable, and mapped to an external framework so the yardstick is not self-defined.

1 · Explainability

Every autonomous action leaves a decision record: what was decided, on what data, what alternatives, which rules, and expected impact — auditable after the fact. *NIST explainability; EU AI Act Arts. 12–13 (record fields are Cognitive Corp's operationalization).*

2 · Human-in-the-loop

High-stakes and irreversible decisions require human approval; routine ones run within defined limits. Boundaries are enforceable, not advisory. *EU AI Act Art. 14; ISO/IEC 42001 A.9.*

3 · Bias mitigation

Space allocation, work-order dispatch, and resource distribution are assessed for equity — bias in building ops has real tenant and staff consequences. *NIST AI RMF Fairness; ISO 42001 cl. 8.4.*

4 · Pre-deployment testing

Before an agent gets write access, it is tested against edge cases, value conflicts, and behavior under pressure — governance testing, not unit testing. *NIST AI RMF Measure/TEVV; ISO 42001 A.6.2.4.*

The fourth capability is the one most often skipped, and the principle behind it is blunt: an agent that cannot demonstrate safe behavior under pressure should not have write access to building systems. Governance testing asks not whether the code runs, but whether the judgment can be trusted — the operational analogue of a safety case, performed *before* the agent is given authority.

How to read this section. The four capabilities above are drawn from external standards, not invented here. Cognitive Corp implements them commercially (as the "Building Constitution" and a pre-deployment testing protocol), and that is a disclosed interest — see p.2. A reader should adopt the *capabilities*, from whatever source, and hold any vendor's claim to provide them against the external frameworks in the Appendix crosswalk.

04 · FROM TOOLS TO CAPACITY

Stop buying tools; start building capacity

The distinction between a tool and a capacity is the practical heart of this report. A tool is bought, installed, and — the buyer hopes — forgotten. A capacity is built, staffed, and improved. Autonomy is available as a tool from dozens of vendors. Governance cannot be bought the same way, because it is not a feature bolted onto an agent; it is the operating discipline that decides which agents are trusted with what, records what they do, keeps humans in the consequential loop, and tests them before and after deployment.

This reframes the buying decision. The right first question is no longer “which autonomous platform should we buy?” but “what is our governance posture, and where are its gaps?” A structured governance-gap assessment — inventory the AI already deployed, score it against the four capabilities and the external frameworks, and map the exposure to the regulatory horizon — turns an abstract risk into a prioritized plan, and routinely reveals that a portfolio’s most consequential autonomous decisions are its least governed, because they arrived quietly through a pilot never designed to scale.

EXHIBIT 4 TWO OPERATING POSTURES FOR THE SAME AUTONOMOUS CAPABILITY

Dimension	Tool-buyer (autonomy only)	Capacity-builder (governed autonomy)
Unit of purchase	A platform / an agent	An operating capability across platforms
Decision record	Vendor logs, if any; not decision-level	Explainable record per autonomous action
High-stakes calls	Automated by default	Human-in-the-loop by enforceable rule
Bias exposure	Unassessed	Audited on space, dispatch, resource decisions
Pre-deployment	Functional testing (does it run?)	Governance testing (can the judgment be trusted?)
Regulatory position	Lower apparent readiness at the deadline	Greater alignment with the deployer duties above
When AI errs	A shrug and a vendor ticket	A record, an override, a correction, an audit trail

Framing based on the governance criteria and regulatory analysis in this report[5–8].

The capacity-builder is not slower to adopt AI. Paradoxically, it can adopt faster, because it can grant autonomy with a safety net rather than a leap of faith. Governance is what lets an organization say yes to the next agent with confidence — which, at a two-year clock speed, is the real competitive advantage.

IMPLICATIONS

What it means

For CRE & workplace leaders

Treat governance as a budget line, not an afterthought. Before the next autonomous platform, run a governance-gap assessment across what you have already deployed. The mispricing is fixable only if you price it.

For FM & building operators

Insist on decision-level explainability and enforceable human-in-the-loop boundaries in every AI procurement. A security certificate is necessary and not sufficient; ask the right-hand column's questions (p.12).

For boards & investors

Ungoverned autonomy is an unquantified operational, regulatory, and litigation exposure. Ask management for the decision record and the regulatory-readiness date. "The vendor is compliant" is not board-grade assurance.

The through-line is a single reframe. The AI-in-buildings market has spent a decade proving that autonomous systems *work*. The next decade will be decided by whether the organizations deploying them can *govern* what they do — explain it, bound it, oversee it, test it, and prove all four to a regulator, an insurer, or a court. The demand is real, the capability is bought, the regulation is dated, and the publicly documented layer that ties them together remains, for now, incomplete. That is the opportunity and the risk, in the same sentence.

Limitations

Public-source basis	"Absent" and "0 of 8" mean no public evidence that a platform documents every applicable governance criterion as of July 2026 — absence of public documentation, not proof that governance does not exist. A platform may govern privately or under NDA.
Point-in-time	Demand, market-size, regulatory, and vendor readings are a mid-2026 snapshot; refresh before reuse, especially the AI Act timeline, amended in 2026 and pending Official-Journal publication.
Attributed figures	JLL and analyst figures are attributed, not independently audited. The market-size figure measures the whole building-automation market (sensors, controls, services), not autonomous-decision AI specifically.
Directional measures	Exhibit 2 positions are group-level, directional readings of public evidence, not item-level vendor rankings.
Not legal advice	Provider-vs-deployer status and high-risk classification are fact-specific, turning on intended purpose and modification; confirm any legal position with qualified counsel.
Conflict of interest	Cognitive Corp sells governance advisory; its framework is self-attested and excluded from all counts (p.2).

Scope	Commercial building operations; the EU AI Act is the only regulatory regime analyzed in depth. Do not extend to other asset classes or jurisdictions without re-scoping.

References

1. JLL, *Future of Work Survey 2026* (published 14 July 2026), based on 2,200+ C-suite and CRE leaders across 21 countries, fielded January–April 2026. Figures cited: 78% expect AI to significantly reshape real-estate strategy; 15% at the “optimization” stage (46% monitoring, 40% analyzing); 60% expect workforce growth. jll.com/insights/future-of-work-survey (accessed July 2026).
2. JLL Research, “Artificial intelligence and its implications for real estate” (accessed July 2026). Figures: 89% of C-suite believe AI can help solve major CRE challenges; 90%+ plan to use AI supporting human experts within 5 years; 60%+ have piloted AI. Macro figures (e.g., projected GDP uplift; US private AI investment) are JLL’s synthesis of third-party sources (incl. the Stanford AI Index) and are attributed, not independent.
3. JLL, “The future of corporate real estate in the AI age” (senior-executive interview program, 40+ executives; accessed July 2026). AI-native business cycles compressing to 12–24 months; system thinker / builder / operator roles.
4. MarketsandMarkets, *Intelligent Building Automation Technologies Market* (edition accessed July 2026): ~$122B (2026) → ~$191B (2030), ~11.8% CAGR. Scope: sensors, controls, connectivity, computing and services — not an AI-agent-specific measure; live figures vary by edition.
5. Regulation (EU) 2024/1689 (EU Artificial Intelligence Act), OJ L, 12 July 2024; in force 1 August 2024. Art. 6 & Annexes I/III (high-risk classification); Art. 13 (transparency), Art. 14 (human oversight), Art. 26 (deployer obligations); Art. 5 prohibited practices (incl. workplace emotion recognition), enforceable 2 February 2025.
6. Council of the European Union, “Artificial Intelligence: Council gives final green light to simplify and streamline rules” (29 June 2026); 2026 Digital Omnibus (political agreement 7 May 2026) deferring high-risk (Annex III) obligations to 2 December 2027 and Annex I product-embedded AI to 2 August 2028; application dates subject to Official-Journal publication.
7. EU AI Act, Art. 99 penalties: up to €35M or 7% of worldwide annual turnover (prohibited practices); €15M or 3% (other obligations, incl. high-risk); €7.5M or 1% (supplying incorrect, incomplete or misleading information).
8. Cognitive Corp, public-source governance review of eight selected smart-building AI platforms (captured June–July 2026). Four criteria mapped to NIST AI RMF, ISO/IEC 42001 and the EU AI Act; each scored cell records platform, criterion, score, source URL and capture date. Named platforms and full evidence register available to qualified reviewers under NDA. Cognitive Corp’s own framework excluded from all counts.
9. Cognitive Corp framework materials (Building Constitution; pre-deployment governance-testing protocol) — self-attested; cited only as one illustrative implementation of externally-defined capabilities (see Disclosure, p.2).
10. European Commission, “Navigating the AI Act” guidance on provider vs deployer obligations (digital-strategy.ec.europa.eu, accessed July 2026); GDPR Art. 22 (automated individual decision-making), per European Commission guidance.
11. ISO/IEC 42001:2023, *Information technology — Artificial intelligence — Management system*. Clause 8 (Operation): 8.2 AI risk assessment, 8.3 AI risk treatment, 8.4 AI system impact assessment; Annex A controls (incl. A.6.2 AI system life cycle / verification & validation; A.9 use of AI systems).
12. NIST, *AI RMF – ISO/IEC 42001 Crosswalk* (airc.nist.gov, accessed July 2026), used to validate the framework mappings in the Appendix.

APPENDIX

Governance yardstick — external-framework crosswalk & sample evidence

The four governance criteria used in this report are not Cognitive Corp definitions. Each maps to a recognized external framework, so the test can be audited against a standard the reader already trusts.

Criterion	NIST AI RMF	ISO/IEC 42001	EU AI Act
Explainability	Explainable & Interpretable (MAP/ MEASURE)	Annex A.6.2 (life-cycle documentation); A.8 (information to interested parties)	Art. 12 (record-keeping); Art. 13 (transparency)
Human oversight	GOVERN 1.x accountability	Cl. 5.3 roles; Annex A.9.2 (responsible use / oversight)	Art. 14 (human oversight)
Bias / fairness	Fairness & harmful-bias management (MEASURE)	Cl. 8.4 (AI system impact assessment); Annex A.7 (data for AI)	Art. 10 (data governance), Art. 15
Pre-deployment testing	MEASURE 2.x test/eval; TEVV	Annex A.6.2.4 (verification & validation); cl. 8.2–8.4	Art. 9 (risk management), Art. 15 (accuracy/ robustness)

Clause- and control-level references are a conceptual mapping, validated against the published NIST AI RMF–ISO/IEC 42001 crosswalk;[12] confirm exact clause numbers with an ISO/IEC 42001 practitioner before any certification use. For the avoidance of doubt, ISO/IEC 42001 Clause 8 comprises 8.2 AI risk assessment, 8.3 AI risk treatment, and 8.4 AI system impact assessment; verification and validation are addressed in Annex A.6.2.4.[11]

Sample evidenced cells (illustrative, anonymized to Platform A–C)

To demonstrate auditability rather than assert it, a sample of scored cells is shown below. Platform identities are anonymized in this public edition; the named platforms and the full scored evidence register are available to qualified reviewers under NDA.

Platform	Criterion	Score	Evidence basis (public, captured Jun–Jul 2026)
Platform A	Explainability	**Absent**	Product & trust pages document security certifications and analytics dashboards; no public evidence of decision-level explanation records for autonomous actions.
Platform A	Human oversight	Partial	Operator dashboards allow manual override of some actions; no published, enforceable boundary policy distinguishing autonomous vs. approval-required decisions.
Platform B	Bias / fairness	**Absent**	No public evidence of fairness assessment for space-allocation or work-order-dispatch decisions.
Platform B	Pre-deployment testing	Partial	Documents model accuracy/QA; no public evidence of governance testing for value conflicts or behavior under pressure.
Platform C	Explainability	**Absent**	Extensive SOC 2 / ISO 27001 material; explainability framed as data privacy, not decision accountability.
Platform C	Human oversight	Partial	Human-configurable rules exist; no published governance framework binding high-stakes decisions to human approval.

How to challenge this appendix

Every “Absent” here is falsifiable: if a platform publishes the missing governance capability, the cell should be re-scored and the finding revised. That is the point of publishing the rubric and the evidence basis rather than a headline number alone. A reader who finds contrary public evidence for a specific cell has, by design, refuted that cell.

Anonymized public edition. “Absent” = no public evidence as of capture date, not proof of non-existence. Cognitive Corp’s own framework is excluded from this review (Disclosure, p.2). *© 2026 Cognitive Corp. May be cited with attribution; methodology is Cognitive Corp intellectual property.*