

AI & CONNECTED INTELLIGENCE

Turn building context into better decisions

How connected building data, analytics and AI move from isolated signals to useful operating intelligence — while keeping what a model concludes separate from what the building is allowed to do about it.

Buildings already generate enormous amounts of data. Most of it does not arrive as intelligence. A thermostat reports temperature. A meter reports power. An occupancy sensor reports a state. A rooftop unit reports mode, runtime and alarms. An air-quality sensor reports CO₂ and humidity. A partner analytic may report a person count or an operational event.

Each observation can be useful by itself. But the building problem is rarely contained in one observation. Power is elevated. Add that the space is unoccupied. Add that the HVAC unit serving it is running. Add that the schedule says the zone should have been in unoccupied mode for three hours. The individual signals have become an operating condition — and only now is there something worth acting on.

Data → context → relationship → condition → decision

AI becomes valuable where the problem requires prediction, classification or reasoning beyond what explicit rules and straightforward analytics can do. But AI is not the starting point. The starting point is a connected building that knows what its data means.

SCOPE

This application explainer owns what intelligence does — not where the workload runs. Its job is to show how normalized building context supports analytics, inference, predictions, recommendations and selected governed actions across real building applications, and to be specific about which of those capabilities exist today. Where the workload executes belongs to Distributed Intelligence. How raw equipment becomes usable building data belongs to Connectivity & Integrations. How an application receives permission to affect physical equipment belongs to Protected Control & Resilience. How edge software and models are deployed and managed as a fleet belongs to Managed Edge.

Distributed Intelligence asks: where should intelligence run?

AI & Connected Intelligence asks: what useful thing does that intelligence do for the building?

FIRST PRINCIPLE

Connected intelligence comes before AI.

The building does not need a model to become more intelligent.

A good schedule is intelligence. A threshold is intelligence. A rule correlating occupancy and equipment state is intelligence. A derived condition is intelligence. A statistical trend is intelligence. An anomaly score is intelligence. A predictive model is intelligence. Computer vision can be intelligence. They differ in complexity, uncertainty and purpose — not in whether they deserve to exist.

A useful application begins with the simplest mechanism capable of answering the operating question. Rules and analytics remain attractive wherever the operating condition is explicit, because they are understandable, inspectable and easy to validate. A model earns its place when the answer is genuinely not directly observable.

Explicit rules	Schedules, thresholds, persistence rules, equipment-state logic.
Connected analytics	Trend comparison, cross-signal correlation, derived operating conditions.
Statistical intelligence	Baselines, deviations, abnormal-use patterns.
Model-based inference	Presence, classification, prediction, anomaly scoring.
Predictive intelligence	Energy forecasting, future demand, maintenance risk.
Partner intelligence	Video analytics, specialized models and third-party applications.

The right form of intelligence should be applied to the right problem. A platform that answers every question with a model is not more advanced than one that knows when a schedule would have done.

THE CONNECTED-INTELLIGENCE LOOP

Observe, contextualize, understand, infer, decide, verify.

The model is only one part of the loop, and usually not the part that determines whether the application is useful.

1 • Observe	Acquire the signals relevant to the operating question.
2 • Contextualize	Associate signals with spaces, equipment, schedules, occupancy and building relationships.
3 • Understand	Evaluate multiple signals together to identify the operating condition.
4 • Infer	Where necessary, use statistical or model-based intelligence to estimate something that cannot be directly observed.
5 • Decide	Produce an insight, alert, recommendation or request for a governed physical action.
6 • Verify	Determine whether the conclusion remained valid and, where an authorized action followed, whether the building actually responded.

1 • Observe

Intelligence is limited by what the building can actually see. An application cannot infer its way around a missing measurement forever. If an energy workflow needs occupancy state and the deployment has no credible occupancy signal, that limitation matters. If an equipment-health model depends on power, runtime and temperature but receives only a binary on/off state, the conclusions it can honestly support are different.

So a capability begins with an admission question: does this building have the data required to support it? Before an intelligence capability is enabled at a site, commissioning should establish that the required sensing, equipment state, data quality and operating context are actually available. Accuracy claims are tied to validated sensing configurations and site admission checks, not to universal marketing promises.

What a site missing an input must never receive is a quietly degraded version of the capability presented under the same product claim. The honest outcomes are visible ones. The site can fail admission for that capability. Or, where the degraded mode was itself validated, it can be admitted at explicitly reduced confidence and a narrowed permitted use. An occupancy capability that ordinarily draws on radar, door state and schedule may remain qualified at lower confidence when schedule data is unavailable, if that configuration was validated. If the radar itself becomes unavailable, the remaining inputs may still support analytics while no longer supporting real-time control. Numerical thresholds and point requirements belong in the applicable capability specification and site-readiness checklist, so they can evolve without changing this explainer.

A capability is only as qualified as the sensing the site actually has — and a site should be told which of the two it is getting.

2 • Contextualize

More data does not fix missing meaning. Register 40217 = 724 is a value. **RTU-3 • supply air temperature • 72.4°F • serving Zone 3** is building information. Add zone occupied, room temperature rising, cooling requested, compressor running and electrical power elevated, and the application has an operating condition rather than a collection of numbers.

Context is also what makes intelligence reusable. The same occupancy state can support HVAC, lighting, cleaning and facility-use workflows. The same equipment state can support diagnostics and energy analysis. The same meter data can support benchmarking, optimization and reporting.

3 • Understand

Many useful conclusions emerge from relationships rather than models.

Space unoccupied + HVAC running + elevated power	Possible unnecessary operation.
Occupied space + CO₂ rising + ventilation inactive	Ventilation condition deserves attention.
Cooling requested + equipment running + room not approaching setpoint	Possible performance problem.

Water flow + no expected use + environmental change

Possible abnormal water condition.

None of these inherently requires AI. They require connected information.

4 • Infer

AI becomes useful when the answer is not directly observable. How much energy is this property likely to use over the next interval? Is the observed occupancy pattern normal for this type of space? Is this energy behaviour unusual enough to deserve attention? Does a combination of observations indicate presence? How many people are using a space? Does video contain an operational event that matters to this workflow?

These questions benefit from statistical models, machine learning or specialized third-party intelligence. The production capability set today covers energy consumption prediction, presence sensing, occupancy and space utilization, and energy anomaly detection. Everything beyond that is either partner-enabled or design-partner and later stage, and the maturity table later in this document says which is which. AI video analytics in particular is partner-enabled rather than a LEXI-native model — native and partner intelligence can share the same building context without pretending they are the same product.

OUTPUT

An inference has to be useful to something.

“We ran a model” is not an application outcome. Every intelligence capability should have a named operating consumer before it has a model.

Energy prediction	Helps an energy application understand likely future consumption.
Presence sensing	Provides an operational presence state to occupancy-aware applications.
Occupancy and utilization intelligence	Helps facility teams understand how space is actually being used.
Energy anomaly detection	Surfaces deviations that deserve investigation.
Video analytics	Generates selected security or operational events through approved partner integrations.
Equipment-health intelligence	Can help identify equipment that deserves maintenance attention before a conventional alarm explains the full problem. A later-stage capability — see the maturity table.

A model without a defined operating consumer is a demonstration. A model with a defined purpose becomes a building capability.

CAPABILITY ABSTRACTION

Applications should consume capabilities, not models.

An energy application should be able to request an energy prediction without knowing the internal endpoint, model family or provider that generated it. An occupancy application should consume presence or occupancy state without being rewritten every time the sensing or inference approach changes.

Building data → intelligence capability → application

The model can improve behind that interface. The application should not need to be reconstructed each time it does.

Abstraction is also a compatibility contract

The interface only holds if its meaning holds. A newer model may improve accuracy, efficiency or implementation behind an existing capability, but a material change in the meaning, units, range, validity or confidence semantics of that output should be treated as a different capability rather than the same one improved. Where semantics differ materially, the platform should create a new capability identity or a versioned transformation rather than silently substituting the old one — and an application should be able to declare the minimum semantic version, latency or confidence it depends on. Applications should be able to rely on what a capability means even while the model behind it improves.

PARTNER INTELLIGENCE

A third-party model should not create a third-party operating model inside the building.

Specialized partners can provide inference without requiring LEXI to recreate every model. The condition is that partner intelligence enters the platform as a defined capability output rather than as an arbitrary proprietary payload that every application learns to parse. The output contract gives native and partner intelligence a common operating shape while leaving the provider- and model-specific implementation behind it.

A normalized intelligence output is expected to carry elements such as these. The exact composition follows the capability and the deployed integration.

Capability identity	Identifies what inference or analytic produced the result.
Building context	Associates the result with the relevant property, space, zone or equipment.
Normalized output state	Provides the operational conclusion the consuming application understands.
Timestamp and validity	States when the result was produced and how long it remains useful.
Confidence or quality	Carries a meaningful confidence or quality measure where the capability exposes one.

Model or provider version	Identifies the intelligence implementation that produced the output.
Provenance	Distinguishes a native LEXI capability from an approved partner capability.

Applications consume the normalized output rather than coupling themselves to the partner model API — the same principle used for native intelligence. The provider or model can evolve without forcing the consuming application to be rebuilt.

Where an intelligence output influences an operational workflow, the event history should preserve enough provenance to determine what capability produced the conclusion, when it was produced, what building context it applied to, and what action or operator response followed. The exact implementation depends on the deployed application. The architectural requirement is traceability, not a claim of a universal immutable AI audit system.

IN PRACTICE

The same context, used differently by each application.

Energy

Meters provide consumption. Equipment state identifies what was running. Occupancy shows whether the space was being used. Schedules show what should have been operating. Connected analytics can identify unexpected overnight use, operation during vacancy, consumption outside normal patterns and behaviour inconsistent with schedule. More advanced intelligence adds energy prediction and anomaly scoring.

Occupancy and facility use

An occupied or unoccupied state can inform HVAC and lighting. Cleaning can be scheduled from actual use. Facility managers can compare designed capacity with observed utilization. The application should receive the operational state it needs rather than automatically exposing every underlying observation.

Equipment intelligence

Commanded state, actual state, temperature, runtime, power, fan speed, valve or damper position, ambient condition and historical behaviour can expose mismatches. More advanced equipment-health and maintenance-risk models remain clearly identified as later-stage capabilities until validated for broader use.

Environmental intelligence

CO₂, temperature, humidity and indoor air quality become more useful when combined with occupancy, ventilation state, equipment operation and time. Some conclusions are simple rules; others benefit from trend analysis or more advanced models.

Video and partner intelligence

The platform can receive an approved partner event, associate it with building location, combine it with building state and route it into an operational workflow — without the receiving application having to become a video platform.

A WORKED EXAMPLE

A room is not cooling correctly.

The thermostat reports 74.8°F. That is data. The building context adds: room occupied, setpoint 71°F, cooling requested, fan running, equipment mode cooling, and room temperature still rising. That is a condition.

Connected analytics can conclude that the room is not responding as expected. Depending on the deployed capabilities, the workflow may generate an exception, a maintenance recommendation, an anomaly score, a request for additional verification, or eventually a more predictive equipment-health assessment.

What it does not do, merely because a model produced an inference, is command arbitrary physical action. If the workflow proposes a control change, that request passes through the authority defined by Protected Control & Resilience.

Building state → connected context → analytic or model → operating insight → governed workflow → verified result

AUTHORITY

AI does not own the equipment.

An inference is a conclusion, not an actuator command. A model concludes that HVAC demand is likely to rise. That is an output. An anomaly model concludes that a unit is behaving unusually. That is an output. A vision model concludes that a space is occupied. That is an output. A predictive model recommends changing an operating strategy. That is an output. None of them inherently grants physical authority.

Underneath that argument is a structural fact: field writes are mediated through LEXI's governed control path rather than being issued directly by an AI model or third-party workload. Models and partner applications receive purpose-bound capabilities rather than protocol credentials for equipment; any write authority is granted explicitly, scoped and bounded rather than acquired as a side effect of network access. A model output therefore enters the physical world only as a request, evaluated against configured authority and operating limits before anything actuates. How that evaluation works, and what happens when it fails, belongs to Protected Control & Resilience.

Better intelligence does not automatically mean more control. The two are separate questions, and they should stay separate as the intelligence improves.

UNCERTAINTY

A model output is not a fact because software produced it.

Rules can be wrong because configuration is wrong. Sensors can be wrong because an input is stale. Models can be wrong because operating conditions differ from the environment they were validated in. A useful application preserves that uncertainty rather than discarding it at the interface.

How confident is the inference?	Confidence should travel with the output, not be dropped when it is consumed.
Are the required inputs current?	A high model confidence does not overcome a stale input, an invalid semantic mapping or a failed actuator verification.
Is the site inside its validated conditions?	A capability validated elsewhere is not thereby validated here.
What happens when confidence drops?	The capability should publish revised confidence, and the workflow should respond to it rather than continue unchanged.
What is the consequence of being wrong?	An uncertain utilization estimate and an uncertain physical-control request do not carry the same risk.

The application should convert confidence into an operating response proportionate to the consequence — an alert, a recommendation, or a request for action — rather than applying one threshold to every use.

VERIFICATION

Intelligence improves when the building can observe what happened next.

Suppose a workflow concludes that a room should transition to an energy-saving condition. Protected control authorizes the request and the equipment receives the command. The connected building can then observe setpoint readback, equipment mode, temperature change, power, runtime, occupancy state or another relevant physical signal.

This matters more than it sounds. A protocol acknowledgement means a command was accepted, not that anything physically happened. A system that treats transmission success as proof of completed control will report results it did not achieve.

Observe → infer → decide → act → observe again

Operational verification	Did the building behave as expected?
Intelligence evaluation	Was the underlying conclusion useful and accurate?

Those questions should not be conflated. An excellent prediction can still be followed by equipment that fails to respond, and a correct physical response can still result from a model that later proves unreliable elsewhere.

When verification fails

A failed verification changes the workflow. If an authorized physical action does not produce the expected response, the application should not continue reasoning as though the requested state was achieved. Depending on the configured workflow, the condition can remain unresolved, generate an operator exception, inhibit further dependent optimization, request additional verification, or return control toward a defined operating strategy. Protected Control & Resilience governs the underlying fallback and recovery authority.

The failure class also matters, and separating the two classes is what makes either of them fixable. A correct inference followed by equipment that did not respond is a control or equipment-response problem. A physical system that responded correctly to a poor recommendation is an intelligence-quality problem. Systems that cannot tell these apart tend to blame the building for the model, or retune the model to compensate for a mechanical fault.

PEOPLE

Automation should remove repetitive decisions, not hide uncertain ones.

Information	Something changed.
Insight	This condition deserves attention.
Recommendation	This is the likely response.
Bounded automation	An approved action occurs inside defined authority.
Human decision	The operator retains judgment where consequence or uncertainty warrants it.

Intelligence is valuable because it improves decisions — not because it removes people from every decision. An application that quietly escalates itself up this ladder as its confidence improves is making a governance change, not a technical one.

PRIVACY

Minimization is an intelligence design choice, not a policy afterthought.

Many intelligent building applications operate on information about people: occupancy, people counts, dwell time, video, location. A connected-intelligence architecture can ask a better question than “what raw data can we collect?” — namely, what information does the application actually need to do its job?

Raw observation → occupied / unoccupied · count = 14 · event detected · space-utilization state

A binary occupied-or-vacant capability can be supplied to an application instead of raw radar measurements. Where possible, the operational state can leave the edge without the raw evidence

necessarily leaving with it, and an application receives only the information required for its approved function and only for the approved duration.

Minimizing the information exposed is usually more valuable than protecting unnecessary information after collecting it — data that was never moved does not have to be defended everywhere it would otherwise have gone.

SCALE

One room gives immediacy. A portfolio gives comparison.

Across sites, intelligence can compare equipment classes, building types, weather conditions, occupancy patterns, energy performance, operating schedules and recurring anomalies. Commissioning evidence from previously qualified sites can accelerate discovery at a new one, and a resource matching a known equipment family can inherit a prior hypothesis without inheriting unconditional qualification.

Two qualifications belong on that paragraph. First, it does not mean every portfolio model becomes an operating model for every property; buildings have different equipment, baselines and conditions, and local applicability has to be validated before anything learned elsewhere is promoted here. Second, portfolio and fleet scale is a later horizon for the platform rather than a description of current single-site deployments. What is true today is that connected data creates a learning surface broader than any one site.

CAPABILITY MATURITY

Be specific about what exists today.

This is the table the rest of the document depends on, so it is worth being precise about what the status column means. Production - qualified deployments means the capability is built and operating in qualified current design-partner deployments. It does not mean the capability is available at every site or in every configuration; availability still depends on successful capability admission. Live · partner-enabled means the capability is exposed to customers today through an approved partner integration rather than a LEXI-native model, and its availability follows that partner integration. Design-partner / later stage means the capability is being developed or validated with specific customers and should not be sold as a general product function.

Capability	Status	Scope note
Energy consumption prediction	Production - qualified deployments	Requires configured meter and context inputs and successful capability admission.
Presence sensing	Production - qualified deployments	Depends on an approved sensing configuration and successful capability admission.
Occupancy and space utilization	Production - qualified deployments	Output quality depends on deployed sensors, site configuration and admission criteria.

Capability	Status	Scope note
Energy anomaly detection	Production - qualified deployments	Requires connected energy and operating context and a suitable baseline.
AI video analytics	Live · partner-enabled	Supplied through an approved partner integration rather than a LEXI-native model; outputs are associated with normalized building context.
Demand forecasting	Design-partner / later stage	Site and data requirements are capability-specific.
Equipment health and maintenance risk	Design-partner / later stage	Requires equipment-specific data, validation and a defined application scope.
Dwell-time analysis	Design-partner / later stage	Depends on sensing configuration, privacy model and capability admission.
Environmental analytics	Design-partner / later stage	Required environmental and context inputs are deployment-specific.
Asset location intelligence	Design-partner / later stage	Depends on compatible location technology and site admission.

Broader data-intelligence work — digital twins, fault detection and diagnostics, operational forecasting, predictive maintenance — describes the direction of the platform rather than a set of capabilities with maturity assigned to them today. Nothing in that domain should be presented to a customer as a dated commitment.

The production capability set is deliberately narrower than the opportunity. That distinction should remain visible in customer and partner materials, because a maturity table is only useful to a buyer who believes it.

REUSE

One intelligent output can support more than one application.

Presence can support HVAC, lighting and occupancy. Occupancy can support facility use, cleaning and environmental interpretation. Energy prediction can support energy-management workflows and reporting. Equipment-state intelligence can support diagnostics and energy. An operational video event can support a security or facility workflow.

The building becomes more intelligent by reusing understanding — not by creating another silo for every model. This is the practical payoff of the capability abstraction described earlier: the second application to need presence does not need a second presence system.

WHAT THIS MEANS

For the people building and operating the system.

Building owners

AI becomes an extension of connected-building operations rather than a separate technology project. Owners can begin with applications that have clear operating value and introduce more advanced intelligence where it improves a defined decision.

Operators

Intelligence should reduce the amount of raw data an operator has to interpret — surfacing what changed, what it likely means, why it matters, what evidence supports it, whether action is required, and whether the building responded afterward.

Systems integrators

Field engineering becomes reusable. Equipment relationships, point semantics, location, control scope and commissioned workflows become context that additional analytics and AI can consume later, rather than work that has to be repeated for each new application.

AI and technology partners

The platform provides equipment connectivity, protocol integration, building context, edge infrastructure, remote lifecycle, data boundaries and governed physical authority — so a partner can focus on the intelligence they actually differentiate through.

CURRENT AVAILABILITY

Connected intelligence is current; the AI capability set expands progressively.

LEXI already uses connected analytics and selected production intelligence capabilities in energy, occupancy and building applications. Capabilities operating in qualified production deployments include energy consumption prediction, presence sensing, occupancy and space utilization, and energy anomaly detection. AI video analytics is also current as a partner-enabled application.

Demand forecasting, equipment-health and maintenance-risk analysis, dwell-time analysis, environmental analytics and asset-location intelligence are design-partner or later-stage capabilities rather than universally available production functions. Portfolio and fleet-scale intelligence is a later horizon.

The intelligence available at a particular property depends on sensing configuration, connected equipment, software release, hardware configuration, model or partner capability, data quality, application scope and successful capability admission. Current project, capability and site-readiness documentation governs the specific functionality and input requirements available in a deployment.

THE CORE IDEA

AI becomes useful when it operates on a building that already understands itself.

Connecting the building creates the data. Normalization gives that data meaning. Connected analytics combine signals into useful conditions. AI estimates what cannot be directly observed. Applications turn those outputs into decisions. Protected control determines what may become physical action. Verification closes the loop against the real building. Distributed Intelligence places the work where it belongs across edge and cloud.

The result is not “AI for buildings” as another isolated software layer. It is a connected building in which increasingly capable intelligence uses existing context, infrastructure and governed workflows to solve actual operating problems.

Connect the data. Give it building context.

Use rules where rules are enough. Use AI where inference adds value.

Keep intelligence separate from authority. Verify against the physical world.

Analytics and AI capabilities vary by application, sensing configuration, hardware, software release, connected equipment, partner integration and deployment. Model outputs are not physical commands. Any physical action remains subject to configured permissions, operating limits and control policy.

Related · Distributed Intelligence · Connectivity & Integrations · Protected Control & Resilience · Managed Edge · Energy, HVAC & Lighting · Equipment Monitoring & Diagnostics · Occupancy, Environment & Facility Use