

EQUIPMENT MONITORING & DIAGNOSTICS

See what equipment is doing before sending someone to find out

How equipment telemetry, native diagnostics, operating history and remote verification combine into one building view — so that an abnormal condition can be investigated, narrowed and often resolved before a technician is dispatched.

Most equipment problems do not begin with a diagnosis. They begin with a complaint. A room is too warm. A rooftop unit seems to be running continuously. A pump is not behaving as expected. A boiler reports an alarm. Energy use rises for no obvious reason. A sensor stops updating.

The traditional response is that someone goes to the building to find out what is happening. Sometimes that is unavoidable. But many service visits begin as information-gathering trips rather than repair trips — someone drives across town to establish whether the equipment is online, whether it is actually running, what mode it is in, what was commanded, what happened afterwards, whether there is a native alarm, whether the condition happened once or repeatedly, and whether it is still happening at all.

Many of those questions can be answered from the building when the relevant information is exposed, connected and in context. This application exists to answer as many of them as the available equipment data supports before the drive rather than after it.

The objective is not to eliminate technicians. It is that when someone does go to the site, they arrive already knowing most of what they would otherwise have driven there to discover.

SCOPE

This explainer owns the equipment-operations application layer: equipment status, native alarms, operating history, alerting, remote verification, triage and service escalation. Connectivity and building context come from Connectivity & Integrations; control authority from Protected Control & Resilience; analytics placement from Distributed Intelligence. Diagnostic depth varies with what the equipment exposes — see Current availability at the end.

OBSERVE

Start with what the equipment already exposes.

Many commercial building assets already hold useful operating information. What is available depends entirely on the equipment: a modern communicating controller may expose dozens of useful points, an older device only a handful, and some equipment is monitor-only by design.

Operating state	On or off · occupied or unoccupied mode · heating · cooling · fan state · pump state.
Measurements	Temperature · pressure · flow · electrical demand · current · voltage · runtime · environmental conditions.
Commands and setpoints	Requested mode · setpoint · schedule state · commanded output.
Native alarms	Manufacturer or controller-generated alarm and diagnostic conditions.
Device health	Communications state · last update · sensor status · reported device-health information.

Equipment does not have to be new to become more visible. Where it already speaks BACnet, Modbus or another supported interface, that information is used directly. Where there is no usable digital interface, sensing or a control interface can be added around the equipment. Where it should not be commanded at all, it stays monitor-only. Wireless Building Automation covers those retrofit paths.

Use the manufacturer's diagnostics when they exist

Many modern controllers and packaged systems already produce their own alarms, fault codes and diagnostic states — sensor fault, communication failure, high-pressure alarm, low-temperature protection, service-required. There is no reason to replace those simply because a monitoring platform is present. The native controller often contains the most authoritative equipment-specific fault information; what it usually lacks is the broader building context around that condition — the room, the zone, the schedule, related sensors, energy state or how the same equipment class behaves elsewhere in the portfolio.

That combination changes priority as much as diagnosis. A fault on one rooftop unit is a device condition. The same fault alongside rising room temperature, high occupancy and no redundant equipment serving that zone is an operating priority. The diagnostic comes from the equipment; the priority comes from the building around it.

UNDERSTAND

A point is useful only when you know what it belongs to.

Register 40217 = 724 → RTU-3 · Supply Air Temperature · 72.4 °F · serving Zone 3

Equipment telemetry is associated with the physical building — property, building, floor, room or zone, equipment, operating point — so an operator investigates the equipment as part of the building rather than as an isolated device. The same applies to fan state, compressor command, valve position, electrical demand, alarm state and operating mode.

A single point rarely explains an equipment problem. Consider a rooftop unit: one point says cooling requested, another says fan running, another says supply-air temperature high, another says zone temperature rising, another shows electrical demand present. Together they describe an operating condition that no one of them describes alone — which is why the application evaluates relationships among connected points rather than treating every reading as an independent alert.

TREND**What happened before the alarm matters as much as the alarm.**

A current value answers what is happening now. History answers how we got here — repeated cycling, extended runtime, temperature drift, conditions appearing only at certain times, recurring communication loss, behaviour changing around occupancy or schedule, and the gap between what was commanded and what was observed.

That turns monitoring from a snapshot into a sequence, and the sequence is usually what makes remote troubleshooting possible:

08:00	Occupied schedule begins.
08:03	Cooling requested.
08:04	Fan reports running.
08:06	Supply-air temperature remains high.
08:20	Zone temperature continues rising.

That is an operating story. It may not identify the final mechanical cause, but it changes the service question from “why is Zone 3 hot?” to “cooling was requested and airflow appears present, but delivered air did not cool as expected” — which is a different conversation with a technician, and often a different truck.

Two questions worth asking at design time

Because diagnostics depend on history, two scoping decisions determine what this application can actually see, and both are worth settling before deployment rather than discovering afterwards.

How long is history kept, and at what resolution? Diagnostic usefulness depends on both retention period and stored resolution. A condition that occurred this morning, one that recurs weekly, and a year-over-year equipment comparison impose different history requirements. Retention and resolution therefore need to be confirmed for the deployed application and release rather than assumed from the presence of trending alone.

How often is each point sampled? Sampling rate bounds what is diagnosable. Short-cycling is invisible in a fifteen-minute trend and obvious in a thirty-second one; a slow temperature drift is the reverse. Sampling should therefore match the condition being observed; unnecessarily high-frequency polling increases field-network traffic and data volume without necessarily improving the diagnosis. Supported sampling intervals vary by device, protocol, Gateway configuration and application.

ALERT**Not every abnormal value deserves an alarm.**

A building with thousands of connected points generates an enormous amount of information, and the operational problem is deciding what deserves attention. Alerts arise from several distinct conditions:

Threshold	A value exceeds an expected operating range.
State	Equipment enters an alarm or fault state.
Duration	A condition remains abnormal for longer than expected.
Communication	A device or a critical point stops reporting.
Relationship	Two or more values no longer make sense together.
Schedule or context	Equipment operates when building context suggests it should not.

The purpose of an alert is not to prove a diagnosis. It is to identify a condition worth investigating.

Alerting should reduce noise, not manufacture it

Alarm fatigue can undermine the value of otherwise useful monitoring. Temperature above threshold may be entirely expected during morning startup. Equipment off may be correct because the space is unoccupied. A sensor not reporting may matter — or may be a device someone intentionally removed.

Context is what separates those cases, and it is why alerting is built on equipment, space, schedule, operating mode and related points rather than on point values alone. Initial deployment should include an alert-tuning period. Thresholds, durations and contextual rules should be validated against actual building operation rather than assumed from equipment specifications alone. Alert commissioning is part of commissioning the application.

Beyond fixed alarms

Not every meaningful problem produces a native fault code. Equipment can be running while its behaviour becomes unusual — runtime increasing over time, operation outside expected schedules, commanded state not matching observed state, temperature response changing, repeated cycling, unusual consumption, or behaviour diverging from similar equipment elsewhere. These are identified through deterministic rules, trend analysis and comparative analytics, and where the deployment supports it, model-based techniques.

Abnormal does not automatically mean broken. Setback recovery, extreme outside conditions, unusual occupancy, an intentional operator change or a technician testing something on site are all legitimate explanations for unusual behaviour. The job of analytics is to improve attention and investigation — not to turn a statistical difference into an unsupported mechanical diagnosis.

VERIFY**Confirm what can be confirmed before dispatching.**

Before a visit is scheduled, a set of questions can be answered from the building: whether the condition still exists, whether the equipment is communicating, whether an alarm has cleared, whether an operating state changed, whether a schedule or setpoint is correct, whether related equipment shows the same condition, whether the equipment responded after an authorized change, and whether the issue looks intermittent or persistent.

Any one of those answers can change the response. Several of them together often change it entirely.

Monitoring does not require control

Observing equipment and holding authority to change it are separate capabilities. This application runs perfectly well in a read-only deployment. Where a deployment also permits an authorized operational adjustment — a supported schedule or setpoint correction — that request travels the governed control path rather than the diagnostic application receiving direct write access to field equipment, and remote verification then observes whether the equipment actually responded. Protected Control & Resilience describes that boundary in full.

RESPOND**The first decision is usually not “what failed?”**

It is what to do next, and the useful output of diagnostics is a better answer to that question:

No action	The condition cleared and appears non-recurring.
Continue monitoring	Not urgent, but worth watching.
Remote correction	An authorized operational setting or schedule is corrected.
Planned service	The equipment needs attention, but not emergency dispatch.
Immediate escalation	The condition warrants rapid on-site response.

Escalate with context

When a visit is warranted, the handoff should carry more than an address: which building, floor, room and zone; which asset; what triggered the alert; whether the problem is still active; when it began and whether it has happened before; what other points changed; whether the equipment itself reported a fault; and what has already been checked or changed remotely.

That context is assembled in the LEXI environment and delivered through operator views and notifications. Where an organisation runs a CMMS or service-management system and wants tickets raised there directly, treat that as a per-deployment integration question rather than an assumed capability, and scope it explicitly — Connectivity & Integrations describes what the integration boundary supports.

Remote diagnostics is not remote repair

Bearings fail. Belts break. Filters clog. Refrigerant systems leak. Valves stick. Electrical components fail. No software architecture changes that reality, and any monitoring product implying otherwise should be treated with suspicion.

The objective is narrower and more achievable: reduce uncertainty before field service begins. A technician who arrives knowing which unit, which alarm, when it started, what was commanded, what the equipment actually did and whether the condition is still present starts considerably further into the diagnostic process than one who arrives with a complaint.

THE ECONOMICS

Three numbers you already have.

Field labour is expensive for reasons beyond the hourly rate. A dispatch also consumes travel time, scheduling effort, access coordination and occupant disruption, and it is wasted almost entirely when the technician arrives without the information or the part needed to finish the job.

Rather than offer an industry percentage, the more useful thing is the arithmetic, because most facilities organisations already hold all three inputs:

Cost per dispatch	Fully loaded — technician time, travel, scheduling, access coordination, and the cost of disruption where the space is revenue-generating.
Dispatch volume	Per site and per portfolio, over a representative period.
Information-gathering share	The proportion of those dispatches that produced a diagnosis rather than a repair. This is the number this application addresses directly.

The complementary measure is first-time fix rate — the share of visits that resolve the issue without a return trip. Remote diagnostics moves it by changing what the technician knows and therefore what they bring. Both figures are ones an operations team can measure before and after, which makes this an application whose value can be demonstrated rather than asserted.

Remote monitoring will not eliminate these costs. Avoiding some unnecessary visits, and making the necessary ones more productive, is the realistic claim — and the right measure of the platform is not how many alerts it generated but whether the information improved the operating response.

AT PORTFOLIO SCALE

One building tells you what happened. A portfolio tells you what is unusual.

Once equipment is represented consistently across buildings, the questions get broader: which buildings generate the most recurring equipment alerts, which equipment classes run the longest, which sites repeatedly lose device communications, which units operate outside expected hours,

which conditions recur after being cleared, and which assets behave differently from otherwise similar ones.

Portfolio comparison does not replace local diagnostics — a technician still fixes one machine in one building. It adds a frame of reference that a single site cannot provide, and it is usually where systemic problems become visible: a control sequence deployed wrong across a region, an equipment model that fails early, a maintenance contractor whose sites keep generating repeat calls.

Equipment context should survive equipment replacement

An application tied only to Device 17 / Register 40217 can lose continuity when that device is replaced. An application tied instead to RTU-3 / Supply Air Temperature allows the field mapping to be updated beneath a stable equipment concept, allowing operating history and application context to remain associated with the asset rather than the old device address. Over the life of a building, that distinction helps keep diagnostic history usable as the physical implementation changes.

ANALYTICS AND AI

Intelligence should improve diagnostics, not decorate them.

More advanced analytics identify patterns that fixed thresholds miss — anomaly detection, equipment-behaviour comparison, condition classification, predictive trends, maintenance-risk assessment, and correlation across several equipment or environmental inputs.

Three boundaries stay firm regardless of how good the model gets. A model output is evidence, not a confirmed mechanical diagnosis. It does not automatically create a work order. And it does not automatically acquire authority to command equipment. Distributed Intelligence explains where these workloads execute; Protected Control & Resilience defines what they are permitted to do.

A WORKED EXAMPLE

A better description of the condition, not an unsupported diagnosis.

Room temperature	77.8 °F
Cooling setpoint	72 °F
Occupancy	Occupied
Cooling request	Active
Fan state	Running
Supply-air temperature	Higher than expected
Runtime	Extended
Native alarm	None

A threshold alert reports: room temperature high. A contextual equipment workflow reports: the zone remains above setpoint despite a sustained cooling request and confirmed fan operation, and supply-air temperature is not responding as expected.

That is still not a claim that the compressor failed. It is a considerably better description of the problem the technician has been sent to investigate — and note that the native alarm is absent, which is exactly the class of problem fixed alarms miss.

HOW THIS DIFFERS

Monitoring becomes more useful as context increases.

	Standalone device monitoring / traditional alarm view	LEXI equipment monitoring
Primary view	Individual device, or an alarm and point list	Equipment in building context
Protocol context	Device-specific or BAS-specific	Normalized across supported systems
Related points	Limited, or available with engineering	Evaluated around equipment and space
Operating history	Device-dependent	Combined with contextual operating state
Native diagnostics	Device-specific	Combined with broader building information
Remote verification	Limited or operator-dependent	Part of the operating workflow
Portfolio view	Usually limited or platform-dependent	Common equipment context across sites
Analytics	Device-specific or rule-based	Rules, trends and progressively advanced analytics
Physical authority	Device-dependent, or BAS control logic	Monitoring stays distinct from governed control

The differentiation is not that LEXI invented equipment alarms. It is combining equipment information, building context, remote operations and multiple applications over one connected platform — so the same integration that supports diagnostics also supports everything else the building needs.

WHAT THIS MEANS

For the people who keep buildings running.

Facility and building operators	See equipment state, alerts and history remotely before deciding what needs attention, and resolve some conditions without a visit at all.
Maintenance teams	Begin troubleshooting with operating history and context rather than with a complaint — and arrive with better information about the likely service requirement.
Systems integrators	Reuse equipment mappings and building context instead of engineering a separate diagnostic stack for every application.

Building owners	Visibility into distributed equipment and service activity across properties, and a measurable basis for judging service performance.
Managed-service providers	Turn connected equipment into a repeatable remote-operations service rather than a collection of isolated site dashboards.

CURRENT AVAILABILITY

Diagnostics expand with what the equipment exposes.

Current capabilities include equipment and device visibility, alerts, operating history, remote operating context, and native equipment information where the connected system exposes it. Diagnostic depth varies materially by equipment: a communicating controller exposing operating modes, alarms and measurements provides a far richer picture than a device with a handful of observable points.

More advanced equipment-health analytics, model-based anomaly detection and maintenance-risk functions are delivered according to application, hardware configuration, available data and software release. Retention periods, sampling configuration and any service-management integration should be confirmed against current product documentation for the deployment. The objective is not to promise a diagnosis from every connected asset — it is to make the equipment information that already exists operationally useful.

THE CORE IDEA

Do not send a technician to discover what the building already knows.

Equipment problems will still require people. Parts will still fail. Mechanical systems will still need maintenance. But a connected building can answer a great many questions before the service visit begins. The difference between an information-gathering visit and an informed service visit can materially change the cost and productivity of the response.

Observe the equipment. Put its data in building context. Use history and related points to understand what happened. Alert when a condition deserves attention. Verify what can be verified remotely. Escalate with the information the technician actually needs.

Equipment visibility, native diagnostics, remote operations, analytics and available points vary by equipment, integration, Gateway configuration, application and software release. Physical actions remain subject to configured permissions, equipment limits and operating policies.

Related · Connectivity & Integrations · Universal Gateway · Managed Edge · Protected Control & Resilience · Distributed Intelligence · Wireless Building Automation