The Changing World of HVAC

Why now is the time to invest in your data

Introduction

The ground is shifting. Does your data know that?

For field service professionals, and everyone orbiting the industry, 2026 has already delivered enough to shake even the most solid foundations. The launch of Nvidia's Rubin architecture, accelerating electrification mandates, and refrigerant phase-downs are underway. A skilled labor pipeline continues to narrow as demand for experienced tradespeople expands in every direction.

These are not distant signals arriving on some future horizon. They are here, on jobsites and in facilities, creating immediate decisions that organizations are not equally equipped to make. And alongside the industry's ever-growing environmental requirements and sustainability challenges, they are widening the chasm between the organizations that are ready to face the future head-on and those still operating in a reactive space.

The organizations navigating this moment most effectively share one thing: they know their assets—not in the general sense of a service log or a spreadsheet, but at the level of actual field performance. They understand what is installed, how it is configured, how it is behaving in its specific environment, and where risk is quietly accumulating. They have built a foundation of asset intelligence, and that foundation is the engine driving every decision they make.

This eBook is about why that foundation matters more right now than it ever has before. What is changing, why it is changing faster than most organizations are prepared for, and what the organizations getting ahead of it are doing differently.

Data-fueled performance and environment predictability are the key advantages in riding rapid waves of change.

No matter where you sit along the lifecycle of a piece of equipment—from the manufacturer to the contractor to the asset owner and operator—having instant access to data that reveals the true performance of your equipment in its actual environment is the engine that should be driving every decision. The question is not whether asset intelligence matters. It is whether your organization has built the infrastructure to act on it before the next wave of change arrives.

01

The shift that is already on the jobsite

When Nvidia unveiled its Rubin architecture, the headlines focused on what the chip could do: faster processing, greater efficiency, a new ceiling for AI performance that would reshape the data centers powering it. Those headlines did not cover what Rubin means for the contractors, asset owners, and technicians keeping those facilities running.

Rubin may fundamentally change what data centers need to stay cool. Nvidia's next-generation Rubin chips are built for hyper efficiency, driving a fundamental shift toward high-temperature, direct-to-chip liquid cooling. That change is arriving on jobsites now, whether the field service industry is ready for it or not.

The instinct might be to read this shift as a simplification of new chips + new cooling technology=older equipment out. But that is not how facility transitions work in practice, and it is not what the people managing these systems are actually experiencing.

Carrier CEO David Gitlin has been direct: the future of data center cooling will require both direct-to-chip and traditional chiller cooling technologies running in parallel. Electrification mandates are accelerating. Refrigerant phase-downs are already underway. These are not changes that can be addressed one at a time. They are simultaneous pressures arriving on the same facilities, often managed by the same teams.

Cooling capacity does not disappear because a facility changes its power source. Asset fleets do not flip to high-efficiency models between maintenance cycles. What changes immediately is the complexity of every decision made about those assets.

The immediate challenge: Transitioning a facility toward direct-to-chip cooling requires knowing every asset on site at a level most records simply do not reach. Not just make, model, and serial number, but configuration, refrigerant type, actual energy consumption, service history, and whether a retrofit has changed how a unit behaves from its original specification.

What is at stake when you do not know your assets

The gap between what manufacturer data describes and what field data reveals is where risk accumulates. Manufacturer data describes what an asset should be. Manuals explain how it should be handled. Field data shows what actually happens on the jobsite. For most organizations, none of that information is connected or kept current. The result is that the decisions with the highest financial stakes are being made on the thinnest information.

Knowing which assets are approaching end of useful life versus which ones have years of service remaining is the difference between a well-timed upgrade and an unnecessary one. In a capital-intensive environment, that difference is not a minor efficiency gain. It is the difference between strategic investment and expensive, unforced waste.

At $80 approximate cost to manually decode a single OEM asset model number, across any real portfolio, that math becomes hard to justify.

A large hotel chain that submitted 445,000 assets for analysis found that approximately 50% were at or beyond the end of useful life. This was invisible in any existing inventory because the data to surface it simply was not there. That finding changed what technicians were dispatched to do, what distributors were stocking, and what replacement conversations were happening.

The data flexibility needed to shift a facility from chiller-based cooling to high-heat liquid cooling is massive. It requires an intricate knowledge of each asset on site in order to effectively invest or divest capital in new, higher-efficiency systems. Asset owners, data engineers, and contractors are at a critical inflection point. The opportunity for strategic cost savings and the risk of significant waste are equally on the table right now.

Right now is the time to know how your current assets are performing. Which models are headed for breakdown? Knowing these answers, the shift to high-efficiency assets will be strategic and scalable.

The highest financial-stakes decisions are often being made on the thinnest information.

02

Adapting to changing facility environments and sustainability requirements

Data centers and other critical infrastructure facilities are making the move toward solar, wind, and other sustainable energy sources. What this does not mean is that the need for cooling capacity will be instantly or completely erased. Nor does it mean that entire facilities will shift their asset fleets to high-efficiency models in one transition. Facility upgrades happen in phases, across mixed asset fleets installed at different points in time and designed to different specifications. Managing that complexity requires knowing what you have at a level of detail that most organizations have not yet built.

The intelligence that defines what gets standardized and what gets customized starts with field data.

Pressure arriving from multiple directions simultaneously

Environmental requirements for critical facilities are converging from regulators, clients and tenants, local communities, and from the economics of energy itself. Each pressure is real. None are going away. And together they are pointing toward the same underlying need: knowing what your assets are actually doing in their environments, not what they were designed to do on paper.

Community resistance to large-scale data center development has become a meaningful force in some markets. Local concerns about water consumption for cooling, the energy demand that requires new substations, and the construction disruption that comes with rapid facility buildout are creating friction that project developers did not previously face. The facilities that continue to earn social license to operate are the ones that can demonstrate, with data, how they are managing their environmental footprint; and that requires the same asset intelligence infrastructure that drives operational performance.

Speed-to-market pressures add another layer. The demand for critical facility capacity is not slowing down, which means the construction and commissioning timelines being demanded today would have been unimaginable a decade ago. Approaches that move complex assembly and pre-certification work off-site—including modular construction methods—are gaining traction precisely because they compress timelines without compressing quality. But even these approaches depend on knowing what is being built and why: the intelligence that defines what gets standardized and what gets customized starts with field data from what has already been deployed.

Predictive data as a competitive advantage in sustainability

Field service organizations that have built a foundation of predictive equipment data have a clear advantage when facility requirements shift. They are not scrambling to understand what they have when a new regulation arrives or a new technology requires a transition. They already know. The question of what to do next can be answered strategically, not reactively.

The critical facilities navigating this moment most effectively are the ones where asset intelligence is already embedded in how decisions get made. They know their refrigerant exposure before a phase-down deadline creates urgency. They know their energy consumption by asset before an electrification mandate requires reporting. They know which units are approaching end of useful life before a transition decision forces a replacement under duress. The gap between organizations that have this data advantage and those that do not is widening with every new requirement that arrives.

Organizations that have built a foundation of predictive equipment data have a clear advantage.

03

The AI reality check

The momentum around AI in field service is real. Predictive maintenance, automated diagnostics, energy optimization, AI-assisted dispatch and scheduling are all compelling use cases, and the technology to support them is advancing rapidly. The organizations paying close attention to this space are right to be paying attention.

But there is a hard truth that is not getting enough airtime in those conversations: AI is only as remarkable as the asset data behind it.

The foundation AI needs to actually perform

A retrieval-augmented generation (RAG) system—the architecture that allows an AI model to draw from your own documentation and records rather than only from general training data—is only as useful as the database it draws from. If that base is incomplete, outdated, or missing entire categories of assets, the outputs reflect that. An AI model working from bad data creates false confidence, which may be worse than no prediction at all.

This is not a future risk to plan around. It is a present reality for any organization that has deployed or is considering deploying AI tools against an asset inventory that has not been verified, decoded, and kept current. The AI will work with what it has, and what it has determines what it can do.

Addressing the data problem has to come first. No amount of new AI capability will compensate for an asset record that does not reflect what is actually in the field. Intelligence built on incomplete information is not intelligence. It is sophisticated guessing.

What good asset data actually enables

When the foundation is right, the AI use cases that the industry is excited about become genuinely achievable. Predictive maintenance works when the system knows an asset's actual configuration, its service history, and how its current performance compares to its design baseline. Automated diagnostics work when the technician on site has immediate access to accurate asset records.

Energy optimization works when consumption data is tied to specific assets in specific configurations, not averaged across a fleet or estimated from manufacturer specs that may no longer reflect how a unit actually behaves after years of service and potential retrofits. Dispatch optimization works when risk is understood at the asset level, not inferred from age or service frequency alone.

The Rubin chip does not care whether a facility has a complete asset inventory. Electrification mandates do not pause while a service team figures out which units still run on a phased-out refrigerant. The prerequisite for all of it is knowing what those assets are, how they are configured, and how they are actually performing today.

An AI model working from bad data creates false confidence.

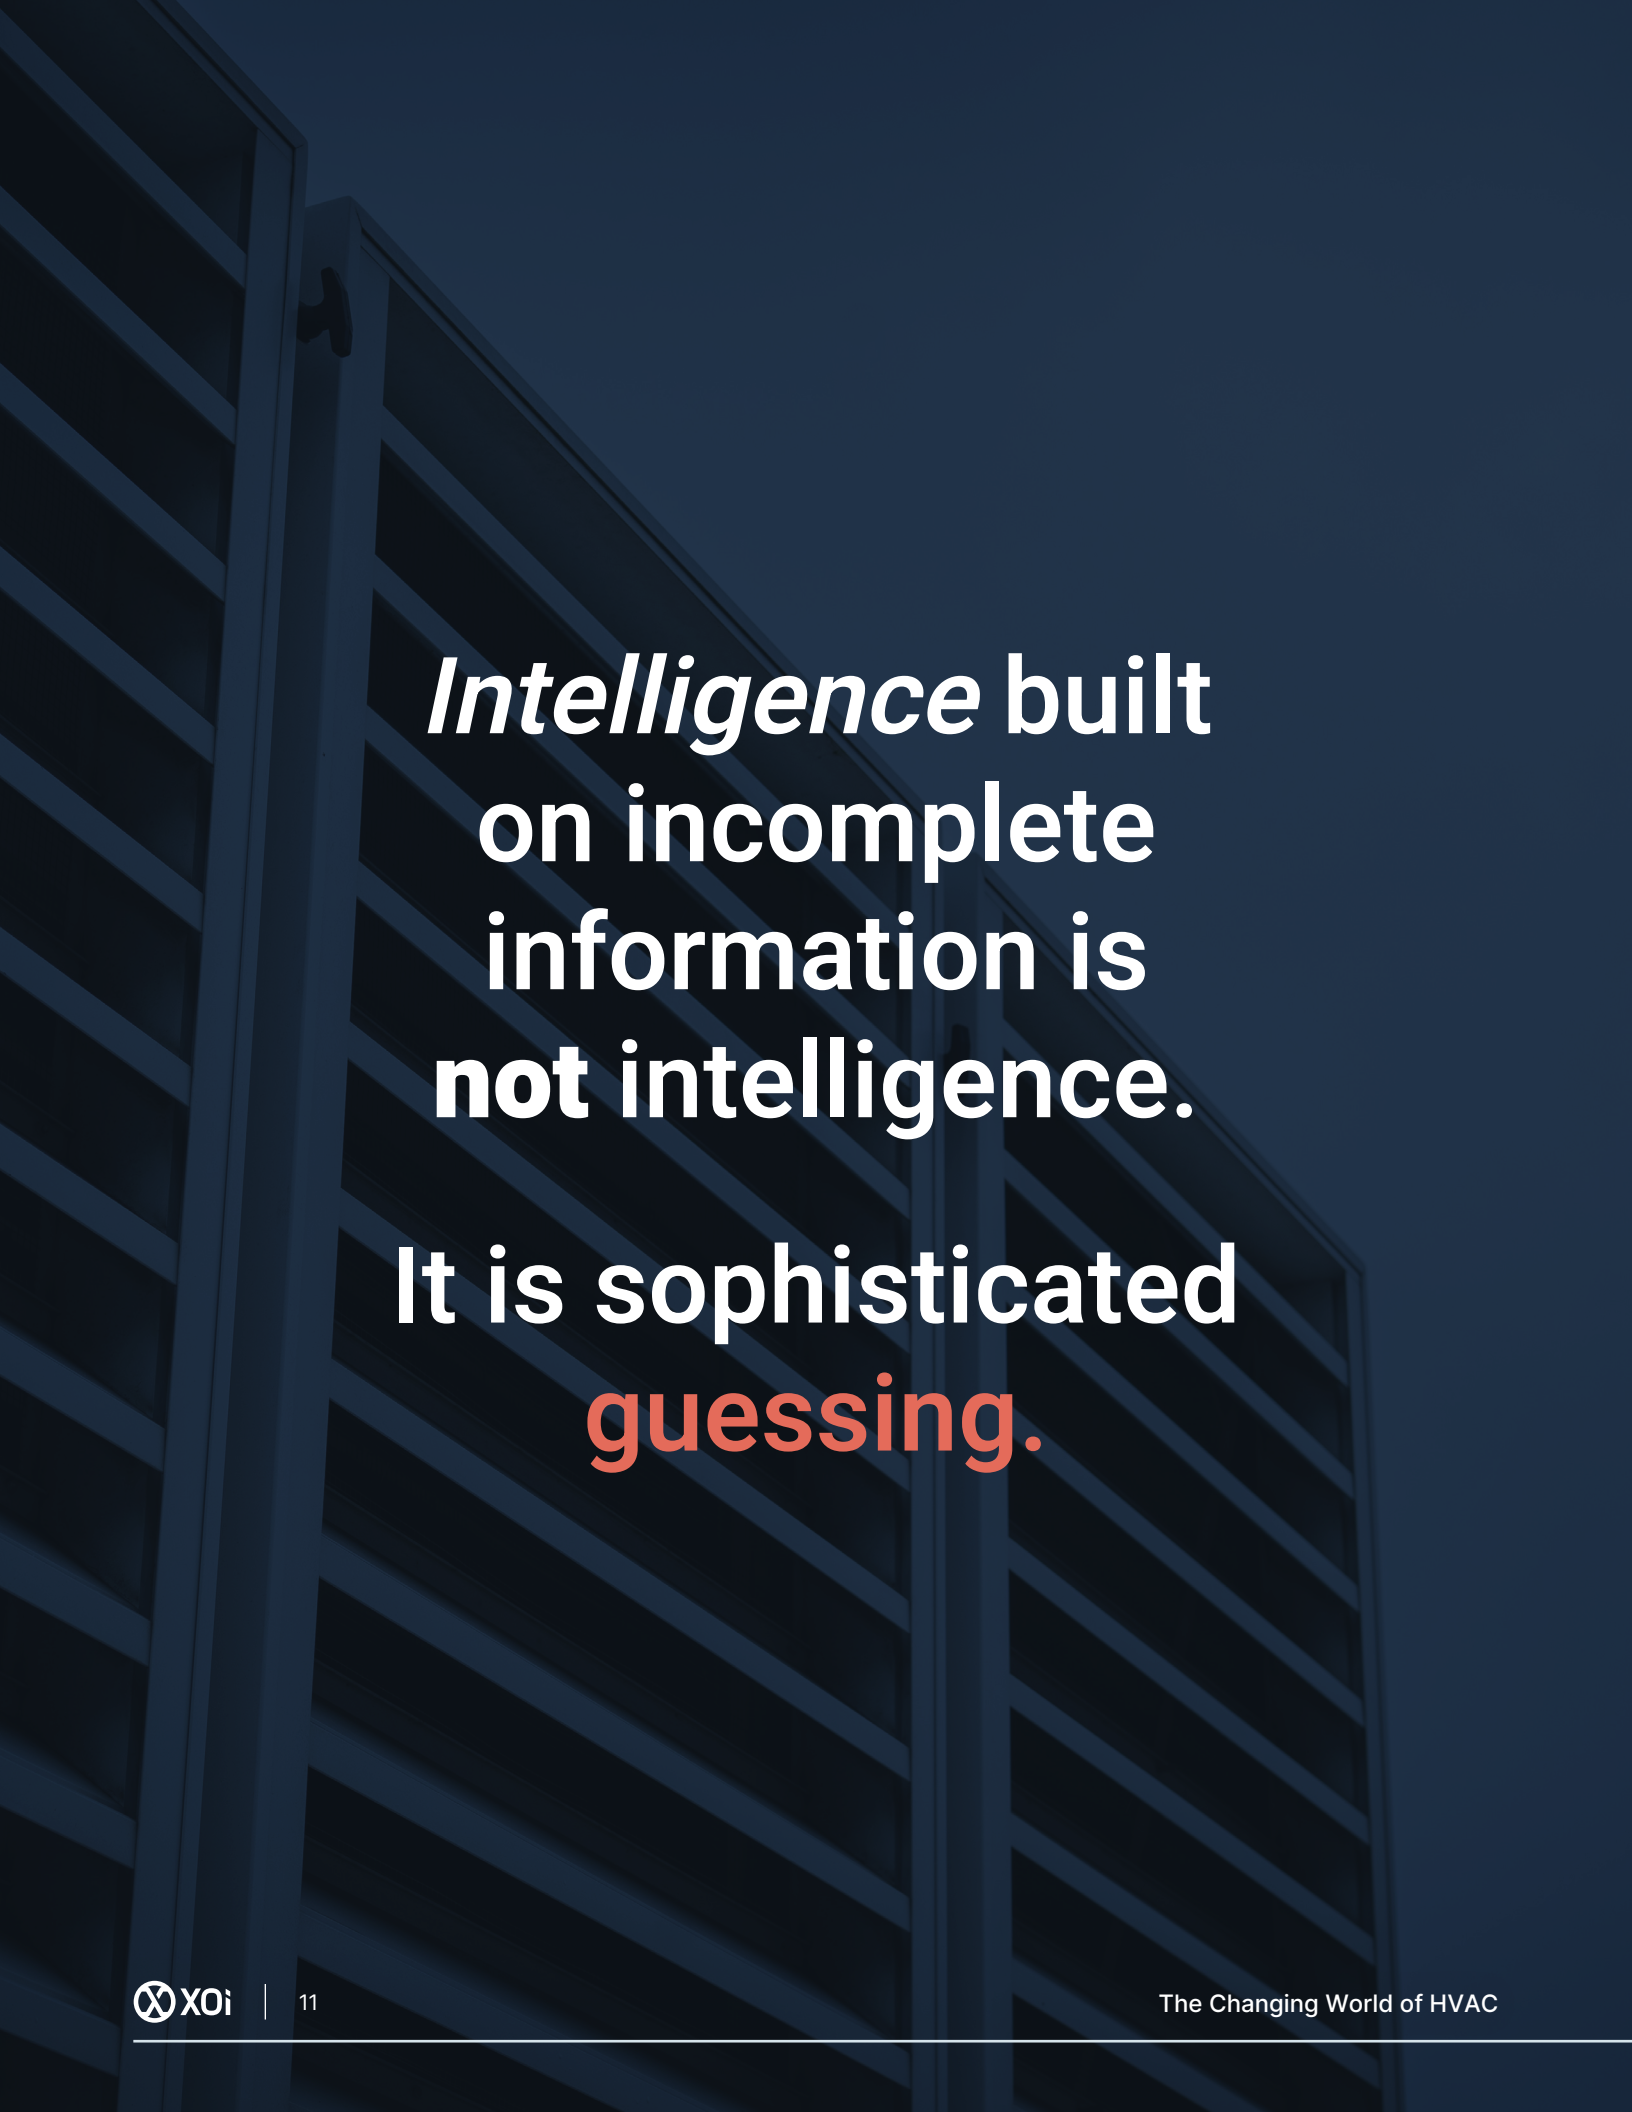

Intelligence built on incomplete information is **not** intelligence.

It is sophisticated guessing.

Why the field technician is the most valuable data source you have

The most accurate, current, and contextually rich asset data in any organization does not live in a database or a manufacturer's spec sheet. It lives in the experience of the field technicians who work on these systems every day. They know which units run hot, which ones have been patched and re-patched, which failure modes appear continuously on a particular make or model in a particular environment.

Historically, that knowledge has had no reliable path from the technician's experience to the systems that make decisions. It lived in conversations, in institutional memory, and then it retired. The organizations building an advantage in asset intelligence right now are the ones creating structured, consistent pathways for field observation to become organizational knowledge. Data captured at the moment of service, routed to the teams that can act on it, and fed back into the decisions that shape what gets built and maintained next.

This is not a small shift. It requires both the technology infrastructure to capture data at the field level and the organizational culture to treat field intelligence as a strategic asset rather than a compliance record. But the organizations that have made this shift are compounding their advantage with every service call. The data from yesterday's job informs the decision on today's. The pattern surfaced across a hundred service visits shapes how the next generation of equipment gets specified.

Every job site is a source of intelligence. The question is whether your organization is set up to capture it, route it, and act on it.

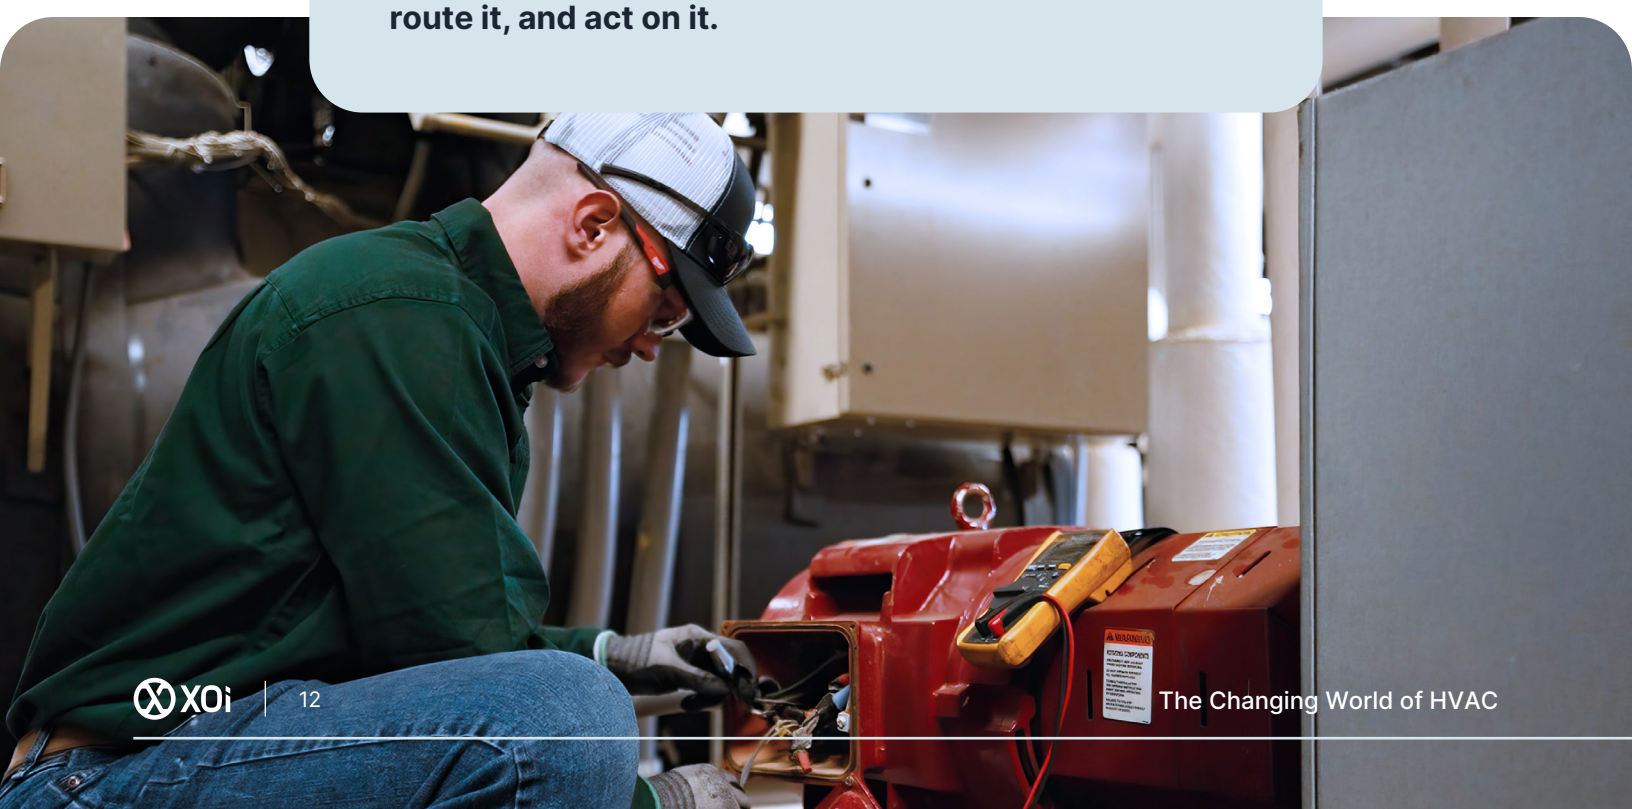

04

The skilled trades gap and what it means for the data you collect

The technology pressures reshaping the field service industry are arriving against a backdrop that makes them even harder to manage. The skilled trades are losing experienced workers at an accelerating rate. Retirements are outpacing new apprenticeships. The hands-on knowledge that walks out the door with a retiring journeyman is not easily replaced.

As you are likely well aware, this is nothing new. What is new is what the convergence of technology transitions, sustainability requirements, and AI opportunity means for organizations that are losing institutional knowledge faster than they are building it.

The connection between workforce and data infrastructure

Here is the dynamic that does not get discussed enough: the skilled labor shortage and the asset data gap are the same problem, expressed differently. The institutional knowledge that experienced technicians carry—which assets are reliable, which are failing, which configurations produce which outcomes—is exactly the intelligence that a well-built asset data infrastructure would capture, preserve, and make available across an entire organization, regardless of who is on the service call.

When that knowledge lives only in people, it's vulnerable. It retires. It transfers to competitors. It is unavailable to the technician dispatched to an unfamiliar site. When that knowledge is captured in a connected asset record and tied to a specific unit, at a specific location, updated with every service visit; it becomes organizational infrastructure rather than individual expertise. It scales, persists, and gets better over time rather than degrading as the workforce turns over.

Approaches like modular construction methods that move complex assembly and pre-certification work off-site, address one dimension of this challenge by reducing the volume of highly skilled work required in the field. But the broader challenge is about whether the knowledge generated by every installation, every service call, and every failure mode gets captured in a form that the next technician, the next project team, and the next generation of AI tools can actually use.

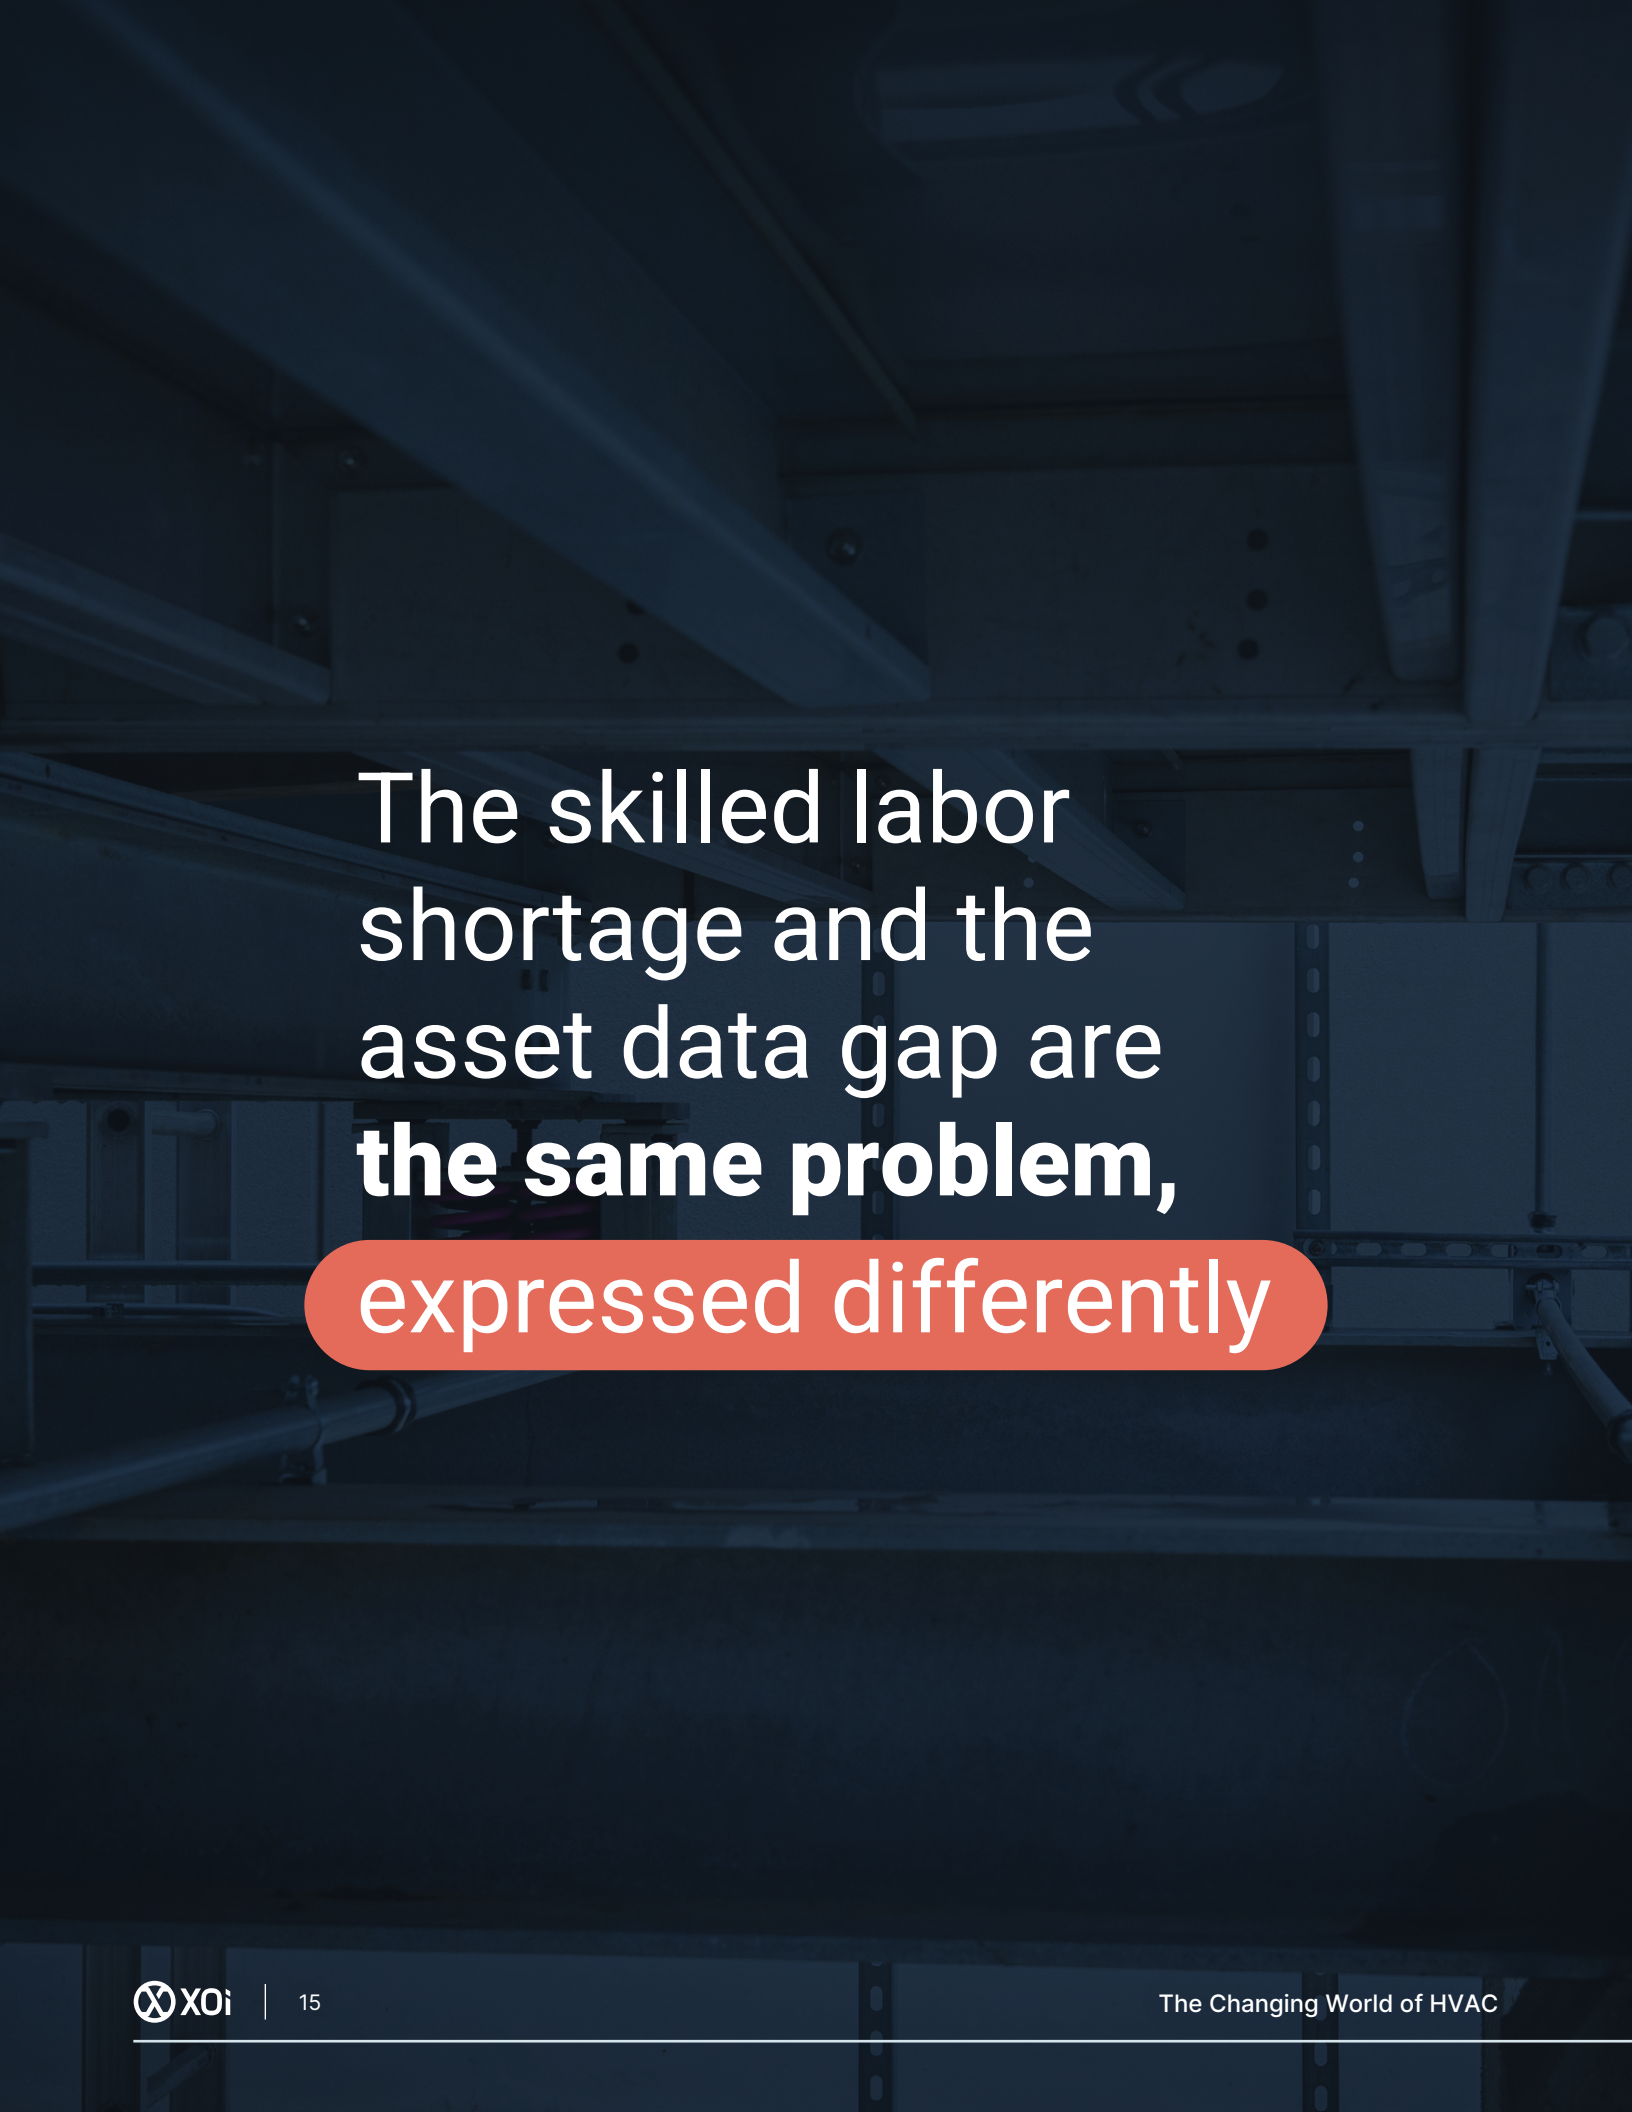
The skilled labor shortage and the asset data gap are **the same problem,** expressed differently

What faster onboarding actually requires

The industry's response to workforce shortages often focuses on training and getting new workers up to speed as quickly as possible. That is necessary. But training a technician to perform a service task is different from equipping them to make good decisions on an unfamiliar asset in an unfamiliar facility.

The latter requires information. What is the asset's service history? What failure modes have appeared on this make and model in this type of environment? What does a unit that is approaching failure actually look like in its diagnostic readings, compared to one that has years of service left? That goes beyond training; it's access to accumulated field intelligence. That access to data is what separates an experienced technician from a new one, regardless of how much classroom time either has logged.

Organizations that have built asset data infrastructure are compressing that gap. A newer technician with access to a complete, current, field-verified asset record can make better decisions faster than an experienced technician working from a spreadsheet and memory. And it has immediate consequences for how organizations manage workforce transitions in a market where experienced labor is scarce.

05

Building the data foundation: what it actually looks like

The case for asset intelligence is not hard to make in the abstract. Most organizations will acknowledge that knowing their assets better would improve their decisions. The harder question is what building that foundation actually requires, and why, despite the obvious value, so many organizations have not done it.

The answer: It's not for lack of trying. It is a combination of where data has historically lived, how it has been captured, and the gap between what a manufacturer's record says about an asset and what the field has actually done to it over years of service.

Where the gap actually lives

Most organizations have some form of asset record that describe the assets as they were configured when they were installed, not as they exist today. A unit that has been retrofitted, had components swapped, been running in conditions outside its design parameters, or accumulated a service history full of workarounds may look perfectly healthy in a static record. In the field, it is a different story.

The gap between the record and the reality is not a documentation failure. It is a structural consequence of how asset information has traditionally been created and maintained. Records are built from installation data. Updates require manual intervention. The technician who observes something significant on a service call has no reliable mechanism to route that observation back into the system that holds the asset record. So the record stays static while the asset changes.

The field technician as the data source

The path to current, accurate, field-verified asset intelligence runs through the people doing the work. Field technicians are the only members of any organization who are regularly in physical contact with the assets.

Building a real data foundation means creating that mechanism. Giving field technicians tools that make capturing structured asset data a natural part of the service workflow, not an additional burden on top of it. Routing what they observe to the dispatch teams, service managers, design teams and procurement functions that can act on it.

When this works well, the intelligence compounds. The observation a technician makes today on a specific asset in a specific facility becomes part of the record for that asset. Pattern recognition across similar assets in similar environments becomes possible. The signal that a particular failure mode is accumulating across a class of equipment surfaces before it becomes a portfolio-level crisis. The decision to replace versus repair, to retrofit versus invest in new equipment, gets grounded in evidence rather than intuition.

The record stays static while the asset changes.

What full-lifecycle data integration actually enables

The most powerful version of asset intelligence is not about any single phase of an asset's life. It is about connecting the data generated across every phase—design, specification, installation, commissioning, service, maintenance, and eventual replacement—into a single, persistent record that travels with the asset.

When that connection exists, the intelligence generated at the end of an asset's life informs the decisions made at the beginning of the next one. The failure modes that appear in service data shape how the next generation of equipment gets specified. The performance data captured in operation tells manufacturers what their designs do in real environments versus what they were intended to do. The service patterns that emerge across a fleet tell asset owners which investments have paid off and which have not.

This is the feedback loop that separates organizations with a genuine competitive advantage.

The organizations that will lead the next decade are the ones treating every job site as a data source, every technician as an intelligence node, and every service call as an input into something better.

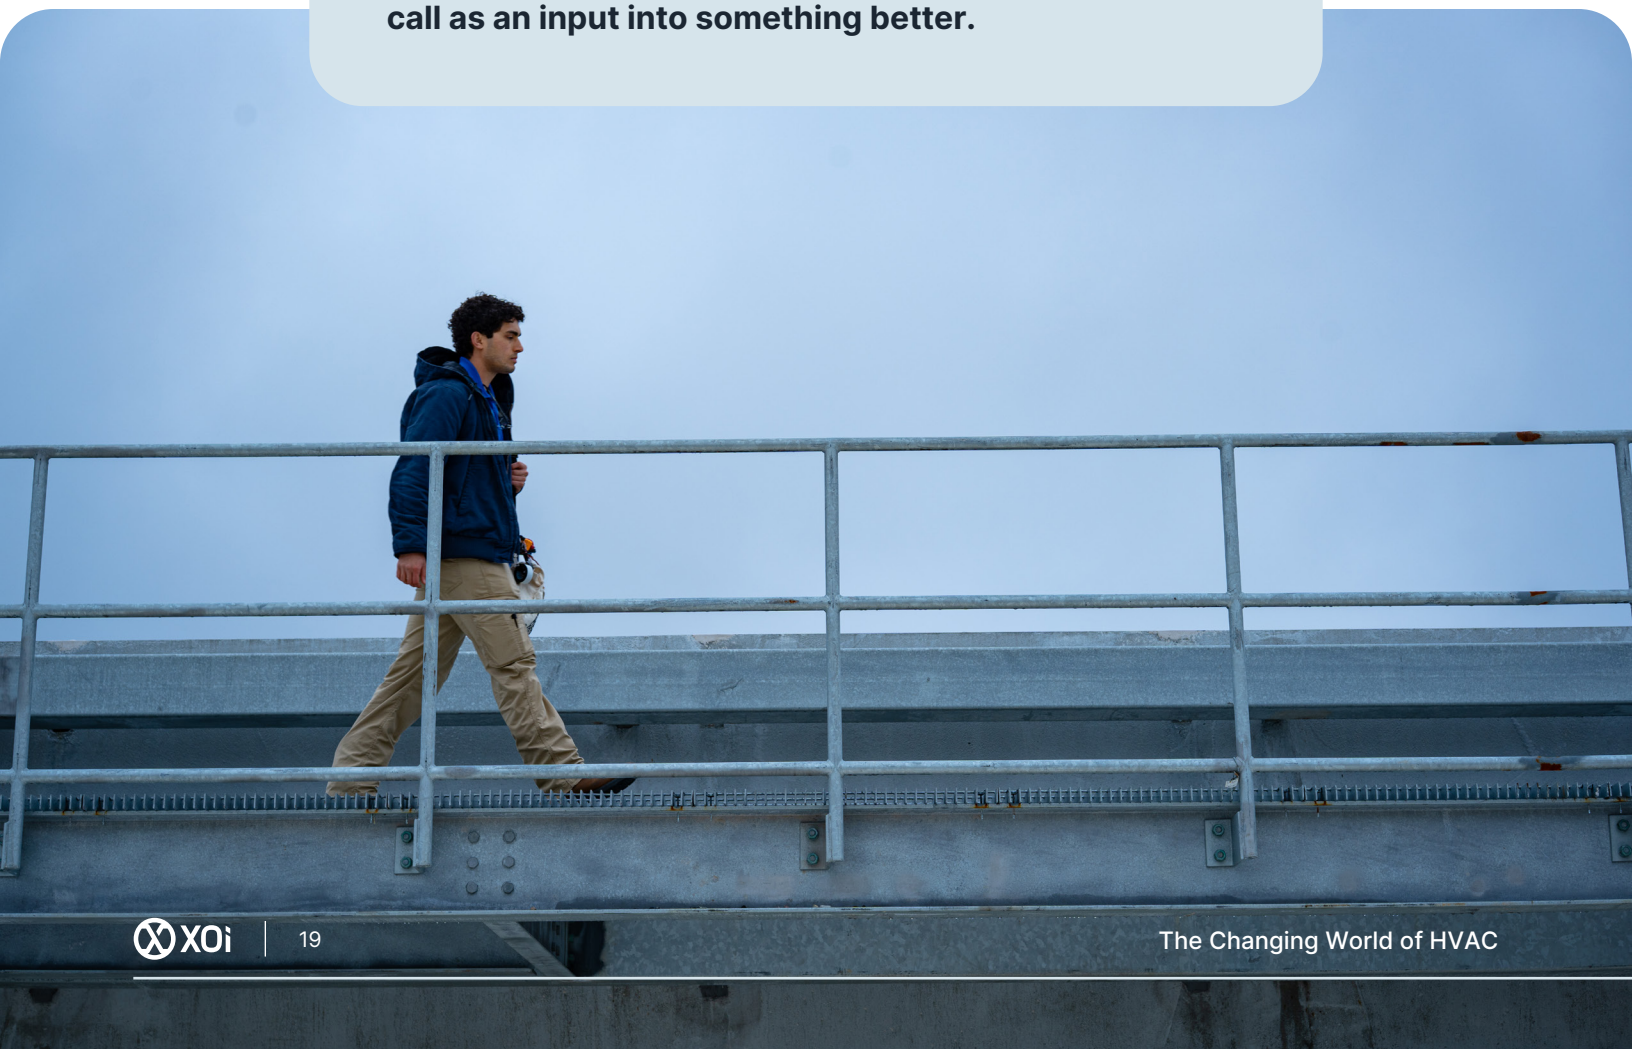

Conclusion

Ready before the next wave. Or waiting for it to arrive.

The announcements of 2026: Nvidia's Rubin architecture, accelerating electrification requirements, and refrigerant phase-downs already underway are magnifying the field service industry's data challenges. The gap between organizations that know their assets and those that do not has always been there. What is new is the cost of standing on the wrong side of it.

The asset owners, contractors, and service organizations that will navigate this moment well are not the ones reacting fastest. They are the ones that built the infrastructure to act before urgency arrived.

That foundation is not built overnight. But it is built incrementally, starting with the decision to treat field data as a strategic asset rather than an administrative record.

With XOi, every job is captured. Every asset is decoded. Every technician observation is routed to the people who can act on it. Over time, the compound effect of every asset recorded and data enriched is an organization that does so much more than just service assets.

XOi helps field service organizations build the asset intelligence foundation they need to compete in a rapidly changing industry.

Learn more at xoi.io