Utilities | Xenon AI

The path to AI implementation in utilities

A practical guide for utilities, from the grid to the customer.

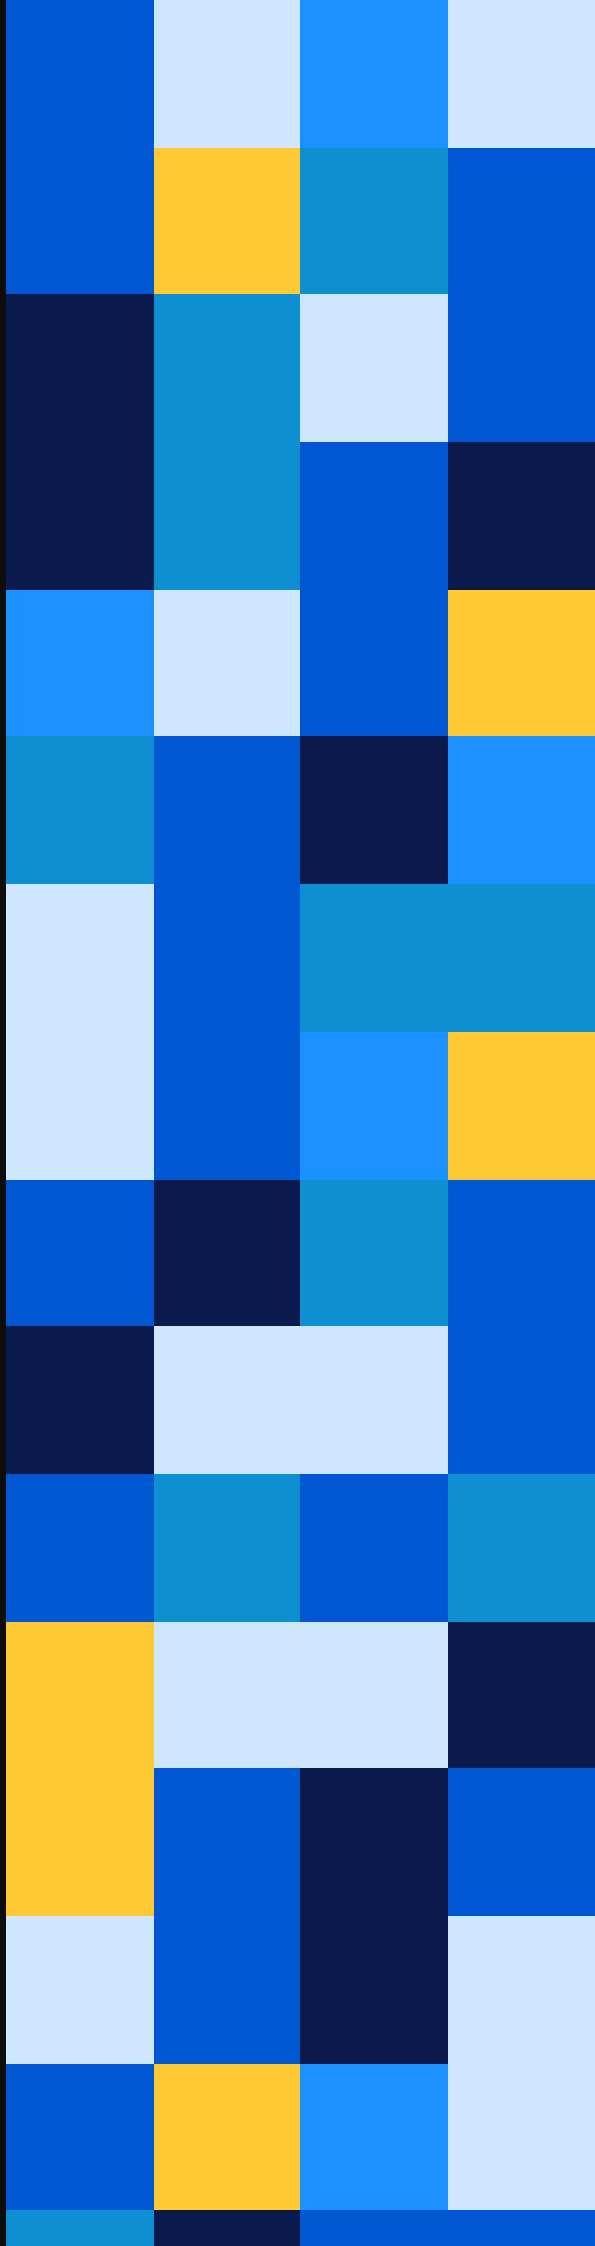

HUDACE

Run Smarter. Grow Faster.

Table of contents

In utilities the pressures are specific: an ageing grid that has to stay reliable, demand that shifts by the hour, rising customer expectations, and a transition to clean energy to plan and fund. AI can help with all of them, but only once you know where to begin. This guide lays out a practical path, from a first honest look at readiness to lasting value, written for utilities from the grid to the customer.

Charting a clear course for AI in utilities

Ask a grid lead, an asset lead, a customer operations lead, and a finance lead where AI should start, and you will hear four different answers: improve reliability, restore outages faster, cut losses, serve customers better. The opportunity runs from the grid to the customer, and so does the temptation to chase all of it at once.

What is usually missing is a route. Deciding to use AI is not the same as knowing which problem to solve first, which grid, asset, and customer data to trust, or how you will recognise a result at the end of a season.

This guide gives utilities leaders that route. It moves through the journey in order, from a candid read of where you stand to the work of holding on to value once a programme is live. None of it is abstract. Each stage reflects how utilities actually run.

Along the way you will see where Hudace and Xenon AI fit across grid, assets, customer operations, sustainability, and finance, so the path stays practical rather than theoretical.

Evaluating your AI readiness

Start with an honest picture of your grid, your customers, and your data.

AI rewards preparation. Before the first model or agent, understand how grid, asset, customer, and finance data flow, and how an ageing network and strict obligations shape what is possible. A grounded readiness check turns interest into progress.

Find your starting point, not a score

Readiness is less about owning the newest sensor and more about the conditions around it: leaders aligned on the goal, data you can rely on from the feeder to the ledger, and teams, in the field and the contact centre, willing to work in new ways. This is not a test to pass. It is a way to see where you are strong and where you still need to build.

A few signs you are ready to take the next step:

- You can tell apart what your people are ready for and what your systems are ready for.
- You can name specific tasks AI could take on: outage prediction, asset and crew planning, load and demand forecasting.
- You know whether grid, asset, and customer data are reachable, accurate, and current.
- You have a real sense of the skills you hold, from the field to data, and the ones to add.
- You can put rough numbers on the time and budget involved.

Done early, this spares you stalled projects later, and lets you scope from facts rather than hope.

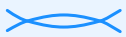

How Hudace helps

Running grid, customer, and finance operations on Hudace means you already have a connected view from the network to the customer, which is a real head start in spotting where AI adds value.

A short readiness session with our team ranks AI opportunities by feeder, asset, and area, so your first projects are the ones most likely to pay off. Talk to Hudace.

Defining strategic AI goals and expected ROI

Tie every AI effort to a number the utility already lives by.

AI earns its place when it moves a number that matters: reliability, restoration time, losses, customer service. Set goals that are specific, owned, and measurable before the work starts.

Clear goals turn effort into outcomes

The most useful projects open with a plain statement of what should change and by how much: faster outage restoration, fewer interruptions, lower losses. Anchor it to a priority, name who owns it, and the work stays focused.

> **The question is rarely whether AI can do the task. It is whether you have decided what a good result looks like, in reliability, in losses, in service, before you start.**

Worth settling early:

- The outcome you are after, written as a number you can track by feeder or area.
- The specific problem, not the broad theme, you are solving.
- A shared view across grid, assets, customer operations, and finance on what is feasible.
- Metrics you are willing to revisit each season.
- A first ROI range, held loosely enough to adjust.

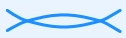

How Hudace helps

Hudace helps you put numbers behind the ambition. Because grid, asset, and customer data already live in the platform, goals and ROI ranges come from what is really happening across your network and areas.

That makes the case for investment far easier to stand behind, and to revisit each season.

20% faster

outage restoration at Gridline Utilities, after connecting grid, customer, and finance operations on one platform. Read the story.

Building your internal AI coalition

Adoption runs through grid, assets, customer operations, and finance alike.

The best model in the world stalls without owners. Progress depends on a small group, drawn from across the utility, who share both the goal and the responsibility for reaching it.

Early on, gather a group that reaches well beyond IT: grid and operations, asset and reliability, customer operations, sustainability, and finance. Their job is not only to comment. It is to own a piece of the change, so it never rests on one team.

This is the group that connects intent to execution. They know which problems are worth solving across the network, and their involvement carries a project past the first season.

What a strong coalition gets right

- It brings the right people in at the start, with a real stake in the outcome.
- It agrees how decisions, risk, and oversight will work before issues arise.
- It leaves room to question, test, and learn out loud.
- It funds the unglamorous parts: enablement, communication, and time.

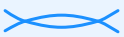

How Hudace helps

If alignment is the hard part, a Hudace discovery session gives your group a structured place to surface use cases across the network and agree on priorities, turning scattered opinions into a shared plan.

When the focus shifts to skills, Hudace Learning offers practical paths so everyone, from the field to finance, feels ready for the change rather than unsettled by it.

Data, the grid, and infrastructure

Good AI depends on good data, from the feeder to the ledger.

AI is only as good as what it runs on. Real-time, trustworthy data, joined across grid, assets, customers, and finance, is what separates a promising pilot from something dependable. In utilities, much of that data also keeps the lights on.

Lay the groundwork for intelligent action

Xenon AI can only reason over what it can reach and trust. That means moving away from data trapped in separate systems toward a connected foundation: grid and asset signals, outage and meter data, customer, and cost, unified and current enough to act on.

Where to focus:

- Data quality: are grid, asset, and customer records clean enough to use without heavy rework?
- Connectivity: can you bring grid, outage, and meter data into one view?
- Infrastructure: can your environment run close to operations and flex with load?
- Ownership: IT keeps systems ready, but grid, customer operations, and finance share it.
- Budget: plan for integration, migration, data quality, and training.

None of this slows you down in the end. It is the difference between AI that demos well and AI you can run a storm on.

How Hudace helps

Hudace gives Xenon AI one governed, real-time view across grid, customer, and finance operations, so prediction, planning, and service work from a single source of truth.

Still untangling older systems? ACE with Hudace shortens the path to a modern, connected core.

Navigating change across the grid and the contact centre

Bring people with you, from the field crew to the service desk.

AI changes the shape of work, not only the tools. The utilities that get the most from it treat the human side as the main event: building skills, adjusting how work is done, and giving people a reason to lean in.

The technology shift rides on a human one

New capability brings honest questions. Will my role change? What happens to the judgement I bring on a switching decision, a crew dispatch, a customer call? Will I keep up? Left unanswered, those questions quietly turn into resistance.

Handled well, this stage is where a controller or an agent stops bracing against AI and starts using it, because it makes their own call sharper.

What helps the shift land:

- Map the skills that are changing and offer real paths to build them.
- Talk early and often, especially where daily work in the field or the contact centre will look different.
- Be straight about changing roles, with AI assisting expertise rather than replacing it.
- Set expectations on pace, season by season.
- Back it with budget for learning, champions, and the culture work that sticks.

How Hudace helps

Hudace Learning gives your teams structured, hands-on paths to grow confident with Xenon AI, from the why through to daily use in the field, in operations, and in customer service.

The result is people who feel ready for the change instead of caught out by it, whatever their role.

Measuring success and scaling AI

A pilot proves the idea. Measurement decides what scales across feeders and areas.

Getting one thing working, on one feeder or one area, is the start, not the finish. The utilities that scale well look hard at what worked and why, then carry that evidence into the next area and the next season.

Let the evidence choose your next move

Useful measurement is not a box-ticking exercise. It is how you learn what really happened, build the confidence to expand, and avoid scaling something for the wrong reasons.

What to track once a pilot lands:

- Measures that reflect real use: reliability indices, restoration time, losses, customer service.
- Actual ROI against what you expected, and the surprises along the way.
- Whether the approach travels to other feeders and areas.
- The resourcing, so people and systems are ready for more.
- What you learned, written down, so the next rollout starts further ahead.

Scaling is not simply doing more. It is doing more of what is proven, with a clear idea of what good looks like.

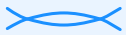

How Hudace helps

Hudace shows you how Xenon AI is used across the business: which feeders, which areas, how often, and to what effect.

That visibility keeps your attention on the work that pays back, and makes the case for the next investment concrete.

Risk, safety, and responsible AI

Value and trust have to grow together, with safety first.

AI does not remove human responsibility. On a live network it raises the stakes on it. Bias, errors, and weak controls are safety, reliability, and trust risks. As AI spreads, the guardrails have to spread with it, and a qualified controller stays in control.

Make trust part of the design

Whether AI is predicting an outage, dispatching a crew, or handling a customer query, the same questions apply: is it secure, is it safe, can you explain the call? Answering them is the job of clear governance, with operations, safety, customer, and IT deciding together how AI is run and watched.

Worth getting right:

- Naming the risks plainly: unsafe switching, biased data, mishandled customer data, unexplained decisions.
- Keeping a qualified controller in control of safety-critical actions.
- Meeting the rules on safety, reliability, and customer protection that apply to you.
- Giving safety, compliance, and model checks a clear owner.
- Treating customer data with the same care as your own.

How Hudace helps

Keeping operations on one platform means less data scattered across systems to defend. Hudace adds granular access controls and built-in compliance at every level.

AI Agent Governance gives you the policies, monitoring, and oversight to keep Xenon AI safe, reliable, and accountable as it grows.

Sustaining value, season by season

Launch is a milestone. Lasting value is the work that follows it.

Going live is the easy thing to celebrate. Keeping value flowing as load, weather, and regulation shift, while you plan the clean energy transition, is the harder, more rewarding work, and it favours utilities that stay curious.

Keep the momentum, and the direction

Maturity does not arrive on launch day. It builds through small iterations, shared learning across feeders and areas, and a willingness to revisit what worked last season. Staying ready for what is next takes both the mindset and the systems to support it.

How to stay ahead:

- Watch how AI performs across feeders and areas, and tune where the numbers point.
- Keep your processes loose enough to adopt what comes next.
- Stay close to grid, asset, and customer teams, and keep learning shared.
- Pair quick wins with the slower investments that make scale possible.
- Keep a habit of small, structured experiments as new options appear.

Lasting value comes from staying adaptable without losing the plot: a more reliable, more efficient, more trusted utility.

How Hudace helps

Hudace helps you keep sight of where Xenon AI earns its keep across the network, so your focus stays on the work that matters.

With Xenon Studio and the wider Xenon AI platform, your teams extend AI at their own pace, and the Hudace Community keeps fresh practice within reach.

Metrics and formulas that matter

AI earns trust when it shows up in numbers you already manage. These are the measures worth instrumenting from the first pilot, with the formulas behind them, so progress is easy to prove and easy to question.

SAIDI

```
SAIDI = total customer minutes interrupted / customers served
```

How long the average customer is without supply.

SAIFI

```
SAIFI = total customer interruptions / customers served
```

How often the average customer loses supply.

Asset availability

```
Availability % = (available time / scheduled time) x 100
```

How much of the time network assets are ready.

Energy losses

```
Loss % = (energy lost / energy delivered) x 100
```

Technical and commercial losses across the network.

Collection efficiency

```
Collection efficiency % = (amount collected / amount billed) x 100
```

How much of what you bill you actually collect.

Recordable incident rate

```
TRIR = (recordable incidents x 200,000) / hours worked
```

The headline read on how safely crews work.

Pick two or three to start. Tie each AI pilot to one, set a baseline before you begin, and review it each season.

Putting Xenon AI to work

A workflow worth starting with, and the questions your teams can ask.

Keeping the lights on, customer by customer: a continuous loop

1. **Sense**
   Xenon AI reads grid, asset, outage, and customer data into one view.

2. **Predict**
   It forecasts load and flags asset failures and outage risks before customers feel them.

3. **Dispatch**
   It sequences maintenance and crews to protect reliability and safety.

4. **Act**
   Owners approve, the platform updates work orders, crew assignments, and customer updates, and the loop learns.

Ask Xenon AI

- "Which assets are most likely to fail and cause an outage, and what maintenance should we schedule?
- "Where are reliability indices worst by feeder and area, and what is driving them?
- "Forecast load by area and time, and flag where the grid is at risk.
- "Show energy losses and collection efficiency by area, and where they are slipping.
- "Where are safety risks rising for crews, and what should we act on first?

Every answer runs on your governed data, so it reflects what is really happening across your operations.

Your AI journey starts at the grid

The next step is closer than it looks.

You do not need every answer to begin. You need a sensible first move, the right people beside you, and support you can lean on. Followed in order, the steps in this guide take a utility from a first honest look to results you can measure, in reliability, in losses, in service.

One feeder or your whole network, the shape is the same: a path that grows with you, where every season teaches you something worth carrying into the next.

Hudace stays with you across that path, from the first readiness conversation to AI working quietly across grid, assets, customer operations, sustainability, and finance, with Xenon AI built into the platform rather than added on.

When your network, your data, and your goals point the same way, the results tend to follow.

Learn more

See AI-native ERP for utilities at hudace.com/industries/utilities.

HUDACE

Run Smarter. Grow Faster.