A BRIEFING PAPER FOR THE ENERGY INDUSTRY · 2026

qurix TECHNOLOGY

Road to AI Excellence

From the status quo to a data-driven, AI-powered energy organization.

With a self-assessment, a maturity model, energy-industry use cases and a 90-day starter plan.

qurix Technology GmbH · Hamburg · qurix.tech

If you are reading this, you are probably thinking about one of these questions.

1. *How do we consolidate our historically grown data landscape — without putting daily operations at risk?*
2. *Where does AI actually pay off? And where is it just an expensive toy?*
3. *How do we finally move from pilot to production — and from production to scale?*
4. *How do I convince the board that this investment will pay back?*
5. *How are comparable companies doing this — and what should we not copy?*

What this document does: It does not give off-the-shelf answers. It gives you mental models, a self-assessment and a 90-day plan so that you can arrive at the answers that fit your organization.

What you'll find on the next pages

You know your data has potential. But where do you start?

Many energy companies know the picture: systems have grown organically. Data sits in silos. AI initiatives fail because the data foundation is missing. And the question "what should we change first?" leads to endless discussions without a decision.

ORGANIZATIONAL

Discussions without decisions

Every meeting ends in consensus that something has to change — but not on what to actually do first.

TECHNICAL

An organically grown landscape

Data spread across ERP, BI, market systems, OT — nobody has the whole picture.

STRATEGIC

AI without a data foundation

Pilot projects burn energy but fail on missing data and governance basics.

Five typical signals that your data landscape is holding you back

If you recognize two or more of these, it's time for a structured step back.

01

“The same number sits in three reports — with three different values.“

Master data isn't consolidated. Decisions are made on shaky ground.

02

“Every AI initiative dies on the data foundation.“

Use cases are clear, but the data isn't in a shape a model can do anything with.

03

“We don't know what market communication will cost us next week.“

Market messaging (MaKo), clearing and billing are a black box, reviewed only every quarter.

04

“None of our pilots ever made it to production.“

There is no handover path from innovation to regular operations. The platform ends at the pilot.

05

“Excel is our most important BI tool.“

Business units build shadow IT because the official self-service tools are too slow or too rigid.

Self-assessment: five questions, ten minutes

Rate your organization on five dimensions — from 1 (not in place at all) to 5 (exemplary). Write down your answers — they are the best foundation for any later conversation about this topic.

Data consolidation	We have a clear, documented single point of truth for our key business objects.	1 2 3 4 5
Data governance	We know who owns which data — including quality, access and deletion.	1 2 3 4 5
Self-service analytics	Business units can build their own analyses without IT tickets — on a consistent base.	1 2 3 4 5
AI in core processes	At least one AI model is in production in a business-critical process (not just a pilot).	1 2 3 4 5
Data culture	Decisions are backed by data — and data work is treated as valuable, not as IT overhead.	1 2 3 4 5

YOUR TOTAL

5–10 Foundations missing — workshop strongly recommended. **11–17** Solid base, clear levers still open. **18–25** Mature organization — workshop focused on the next growth stage.

Maturity model: five stages of a data-driven energy organization

01

Ad hoc

Data everywhere, nobody responsible

Excel and shadow-IT reporting dominates. Data quality cannot be verified.

02

Aware

Pain recognized, first initiatives

First data platform initiative is running. Governance basics are being discussed.

03

Consolidated

Platform in place, master data clear

Single point of truth for the key business objects. BI broadly available.

04

Active

AI in pilot, self-service broad

First AI models in production. Business units build their own data products.

05

Data-native

AI in core processes, a learning organization

Data and models are part of the business model. Decisions are data-based.

A modern data and AI landscape has five layers

Each layer has its own tools, its own owners and its own investment logic. Skipping a layer means building on sand — attacking the top layer too early means building for show.

LAYER 5	**Experience**	*Self-service, dashboards, apps, conversational UI*	Where data and AI work meets the user. Business units have to be able to help themselves.
LAYER 4	**Intelligence**	*AI and ML models, agents, optimizers*	Where the real value addition from data happens — once the lower layers are solid.
LAYER 3	**Modeling**	*Data models, data products, semantic layer*	Where raw data turns into business meaning. Versioned, tested, documented.
LAYER 2	**Integration**	*ETL, market-, ERP-, OT-connectors, master data*	Where data enters the house — reliable, observable, governable.
LAYER 1	**Foundation**	*Lakehouse, storage, compute, identities, network*	The foundation. If this doesn't carry, nothing does.

Where to invest? A prioritization matrix

Four quadrants every initiative falls into. Use this matrix as a first filter before you commit budget.

QUICK WINS

Fast impact, low effort

Consolidate standard reports · data catalog · KPI glossary · first market-communication automation · self-service templates.

STRATEGIC

High impact, high effort

Lakehouse migration · master data consolidation · AI platform · trading signal engine · predictive maintenance for the grid.

BE CAREFUL

Low effort, low impact

Pretty dashboards on weak data · AI demos with no follow-up plan · trainings without tooling.

TRAPS

High effort, low impact

Buying an AI platform without a data foundation · running eight pilots in parallel · big-bang platform migration without a use case.

Five levers that really work in the energy industry

Use cases with high maturity and high value contribution — suitable as an entry or scaling case in a pragmatic roadmap.

FORECASTING	**Load and generation forecasting**	Better forecast quality lowers balancing costs and improves trading margins — proven, clearly measurable.
MARKET COMMUNICATION	**Automating MaKo / market messaging**	Structured handling of the 1,000+ message types reduces clearing cases and clearing-case cost significantly.
TRADING	**Signals from market data**	AI-driven pattern recognition in spot, forward and cross-commodity data delivers trading signals with clear back-testing.
ASSET	**Predictive maintenance for grids and generation**	Sensor data plus maintenance history enable maintenance at the right time — less downtime, less cost.
PORTFOLIO	**Portfolio and flexibility optimization**	Optimal marketing of PV, wind, storage and controllable loads under changing market conditions.

Three typical dead ends — and how to avoid them

01

“Let's buy the AI platform first.“

DEAD END

Tool-driven strategy. The platform is in place — but nobody can fill it with anything meaningful.

BETTER WAY

Use case first. Data foundation second. The tool follows from both.

02

“Let's run eight pilots in parallel.“

DEAD END

Attention scatters, no pilot matures, not a single one makes it to production.

BETTER WAY

At most two parallel tracks. Clear handover criteria into regular operations from day one.

03

“We're waiting for clean data.“

DEAD END

Data is never perfect. While you wait, others learn to work with imperfect data.

BETTER WAY

Define “good enough“. Start with a narrow use case and improve data quality iteratively.

Building the business case: a simple four-quadrant framework

A data or AI initiative pays back through four levers. If you can name all four, you have the business case — if you only have one, think again.

Reduced cost

Manual effort goes down · fewer clearing cases · optimized maintenance cycles · license consolidation.

Example: fewer MaKo clearing cases = lower headcount cost.

Avoided risk

Reporting quality assured · regulatory compliance · audit-readiness · data dependencies transparent.

Example: consolidated master data = less reporting risk.

New opportunities

Trading margin · new market products · demand response · platform business.

Example: flexibility marketing only scales with good data.

Scaling effects

What you build once carries the next use case · data products become an asset · learning curve falls.

Example: the master-data layer will carry the next 5 use cases.

A 90-day starter plan — as a template for your organization

Four phases, clearly separated, each with a single clear deliverable. This is a template — not a claim to be complete.

WEEK 1–2

Inventory

Inventory systems · sketch data flows · collect pain points · identify stakeholders.

→ Heat map of the as-is state

WEEK 3–6

Align the target

Sketch the target architecture · use-case long list · maturity-level placement · internal workshop.

→ Target picture + long list

WEEK 7–10

First quick win

Take one narrow use case to production · measure data quality · share lessons within the team.

→ First visible win

WEEK 11–12

Finalize roadmap

Prioritization · investment estimate · responsibilities · decision paper for the board.

→ Roadmap & business case

The Road-to-Data-Excellence method in five steps

If you would rather not run the 90-day plan alone — this is how we walk it with you.

1

30 min

Discovery call

We understand your starting point, your goals and your constraints.

2

1–2 weeks

Preparation

We review documentation on systems, data flows and pain points and prepare the workshop.

3

1–2 days

Workshop

Together with your team we develop the inventory, the target picture and the roadmap.

4

1–2 weeks

Documentation

You receive a structured document with target picture, roadmap, quick wins and estimates.

5

+4 weeks

Follow-up

We check progress together and adjust roadmap and effort estimates if needed.

What you walk away with at the end of the workshop

Not an abstract strategy paper. A workable concept your team can keep building on the Monday after.

01

Inventory of the as-is

Which systems, data flows and processes exist today? Where are the biggest gaps and redundancies?

02

Target picture — data and AI

What does a modern, scalable architecture for your requirements look like — from data consolidation through AI agents to automated market communication?

03

Concept and roadmap

Which steps move you from as-is to target? What are quick wins, what are strategic investments — with owners and timelines.

04

Business case

What is the value of the implementation — in reduced cost, avoided risk and new opportunities through AI automation.

qurix Technology — AI integration for the energy industry

We are a software-development and consulting firm with a clear industry focus. We build AI solutions until they run in production — speaking the language of the energy industry along the way.

FOCUS

Energy industry

Utilities, grid operators, municipal utilities, generators, trading. We know market communication, EnWG, BSI/KRITIS, the market role model and the day-to-day of utility operations.

PROFILE

Domain plus engineering

Deep energy-industry knowledge combined with modern data and AI engineering. We build all the way into the control system.

REFERENCE

Uniper SE and others

Data platform migration at Uniper SE: +84% performance, −76% platform cost, zero critical incidents. More references at qurix.tech/showroom.

Self-check — fill in or print out

Rate each statement honestly. For scores below 3 a structured look is worth it — for scores below 2 action is overdue.

We know our most important data sources and their owners.	1	2	3	4	5	Reports across the company agree — same number, same understanding.	1	2	3	4	5
Master data (assets, accounts, contracts) is consolidated.	1	2	3	4	5	Business units can build their own analyses.	1	2	3	4	5
At least one AI application is in regular operations (not a pilot).	1	2	3	4	5	We have a data strategy that board and IT share.	1	2	3	4	5
We know our critical market-communication clearing cases and their cost.	1	2	3	4	5	We know which quick wins the next 90 days would deliver.	1	2	3	4	5

Three takeaways we'd like to leave you with

I

Start small — but with architecture sense.

Every quick win should be built so the next one fits onto it. The first use case doesn't have to be big, but it has to point toward the target architecture.

II

Use case before platform.

Tools follow requirements. Buying an AI platform before you have a clear use case costs money and rarely delivers impact.

III

Data governance is a people question, not a tools question.

No tool in the world solves data problems if nobody owns them. Clarify roles first — tools second.

Let's talk briefly.

30 minutes, no commitment. We'll figure out together whether and how Road to Data Excellence fits your situation — and share what we have learned in comparable energy-industry projects.

HOW TO REACH US

contact@qurix.tech

qurix.tech · Colonnaden 18, 20354 Hamburg

Request the workshop directly: qurix.tech/road-to-data-excellence