AI agents inside the energy company

What they can read, what they can change, and who decides

IN BRIEF

01 Energy companies and utilities have started to put AI agents in front of customers and beside their control rooms. A British supplier's agent answered around 8,000 customer emails a week in a trial. An American water district runs a chat assistant and two after-hours voice assistants. An oil major is adding agents that find the cause of equipment faults across its assets.

02 The suppliers are moving in different directions. One control system supplier says its platform can "proactively act on behalf of the operator", and an executive calls it "an on-prem solution". Another supplier's billing agent lists actions to update billing accounts and to stop service, with no word on who approves them.

03 This paper is about the work around the grid rather than on it: the reading, checking and drafting done in field operations, assets, safety, customer service, regulation and engineering. It sets out what an agent can read, how it reaches an energy company's systems, what it may change without a person, and what the rules already say about critical infrastructure, customers and markets. It ends with seven questions to put to any supplier, and with a place to start.

04 The position throughout is simple. An agent prepares. A named person decides. Control systems are read, never written. The agent runs on infrastructure the company controls.

IN THIS PAPER

01 Where energy companies and utilities are today 2
02 The work nobody has packaged 3
03 Three doors, one gate 4
04 What an agent may do 6
05 What the rules already ask 6
06 A worked example: one regional utility 8
07 Seven questions before an agent touches the network 9
08 Where to start 9

PART 01

Where energy companies and utilities are today

Around 8,000
emails a week one supplier's agent answered in a trial

4%
of all the customer emails that supplier received in the UK

Over 13,000
pieces of equipment in an oil major's predictive maintenance programme

Six months
of Katonic's pilot at a state electricity utility, entirely on-premises

Customers meet agents first. Octopus Energy says its agent Arlo "handled around 8,000 emails a week – 4% of all customer emails received by the energy supplier in the UK" during a three-month trial. The company sets limits on it: "every AI-written email is clearly labelled and customers can ask to speak to a human advisor at any stage", and "It never handles vulnerable customers, sensitive cases or complex complaints" [1]. A water district in South Carolina runs a chat assistant, an after-hours voice assistant and a Spanish-language voice assistant, and describes the work as "not a replacement strategy" [2].

Assets come next. Shell will add "AI agent-based root cause analysis and remediation across its asset operations on Microsoft Azure" to its predictive maintenance programme, which already covers "more than 13,000 pieces of equipment" [3].

The control room is the contested ground. Honeywell launched a control room assistant in March 2026 to "proactively advise operators of impending issues", after pilots with Chevron and TotalEnergies [4]. In June it said its Experion Cognition platform can "proactively act on behalf of the operator to help resolve anomalies in the control room" [5]. A trade report quotes the same supplier: "Every situation is verified by your process experts before being available to operators," and the system "allows closed-loop, autonomous action, such as automatically hot-swapping a drifted sensor value" [6].

On-premises agents exist. The same report quotes a Honeywell executive: "Experion Cognition does not require any cloud connectivity. It's an on-prem solution" [6]. IBM's documentation for the AI service behind its asset management assistant explains how to "Deploy watsonx®.ai on-premises in full version mode" [7]. So where an agent runs is not the only question. The harder ones are which tools it may call, as whom, with whose approval, and what record is kept.

What the published accounts leave out. None of the deployments we read describes a step where a named person approves each change an agent makes to an account or an asset record. We found no public incident at an energy company caused by an AI agent, and we do not claim one.

One we can speak to. An Australian state electricity utility, critical national infrastructure where data cannot leave the perimeter. Katonic won the competitive pilot with Cisco: six months, entirely on-premises on a locally hosted model. Reference available on request.

A buyer is being asked to decide how far an agent acts, often before anyone has asked where the person sits.

PART 02

The work nobody has packaged

Every

library template we opened acts or alerts as soon as the workflow decides

None

of them shows a step where a person approves before a ticket, a message or a change to an account goes out

We listed the jobs done around an energy company, from the control room log to the complaints desk, and looked for each in the template libraries of the agent and workflow platforms. The libraries cover the alert, the bill query and the dashboard well. They thin out wherever a named person signs something.

WHERE TEMPLATES EXIST	WHERE NONE WAS FOUND
Smart meter anomaly alerts	**A permit to work checked before the permit authority signs**
Alarm questions in the control room	**A switching schedule checked against the procedure**
Bill and account queries	**A control room shift handover, read only**
High bill analysis and draft replies	**An inspection or integrity record review with engineer sign-off**
Trading market summaries	**A regulatory return assembled for the manager to submit**
Job summaries for field technicians	**Hazard study and change actions with owner and date**
	A complaint response a handler approves
	A contractor's submittal checked against the specification

Two findings matter more than the list. First, every library template we opened acts or alerts as soon as the workflow decides. One smart meter template "Routes critical anomalies through a short cooldown wait and marks them ready for critical dispatch", with no review step on the page. One supplier's billing agent lists "Create & Update Billing Account In Billing System", "Submit Metre Read" and "Stop Service" among its actions, with no approval language on the page [8]. None shows a step where a person approves before a ticket, a message or a change to an account goes out [9].

Second, the templates we opened are built from hosted services: a hosted model, a hosted chat or messaging service, a hosted ticket or billing system. That settles where customer and grid data go before anyone has asked.

The gaps sit where the reading is slow, the records are controlled, and the decision already belongs to a named person. That makes them the safer place to begin.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In an energy company that means three kinds of thing.

Documents.

Operating procedures, safety rules, engineering standards, licence conditions, tariffs, the complaints procedure. These are indexed once as knowledge sources, with each reader's permissions kept, so an agent answers with the document it relied on and says when the documents are silent.

Systems of record.

Work orders, the asset register, inspection records, customer accounts, meter reads, outages. Most are reached through an interface the company already runs or its supplier publishes. The company gives the platform the description of that interface, chooses which operations become tools, and decides which of them need approval.

Grid and plant data for analysis.

Loading, faults, test results, meter data and costs, held in the historian's copy, the data warehouse or a reporting database, and read only. The operational network is not one of the doors.

All three pass one gate: every tool call is checked against the company's policy, which can allow it, record it, hold it for a person, or block it.

Where the company's interface checks the caller, the agent acts as the person using it and cannot claim another.

SIX WAYS INTO AN ENERGY COMPANY'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	A permit, a switching schedule, a customer's letter, an inspection report	Nothing is connected. The safest start
02	**A knowledge source**	Procedures, safety rules, standards, licence conditions, tariffs	Permissions are kept per reader
03	**An imported description of an interface**	Work orders, assets, customer accounts, meter reads, outages	You choose the operations. Changes wait for approval
04	**A reporting database**	Grid and plant history, meter data, costs	Read only
05	**A server from the tool catalogue, run in your cluster**	Databases, web search, document conversion	It runs inside; check what each tool does
06	**Public market and system data**	Prices, system margins, generation, published outages	Reads only, from the publisher

PART 03 · CONTINUED

19
interface descriptions taken from publishers and standards groups; the importer refused none

1,197
operations the importer found in them

754
of those operations read; 443 are sent as changes

326
operations in the British balancing market's data service, and every one reads

What is published. We took 19 interface descriptions from publishers and standards groups: market, grid, weather and carbon data, and the standards for demand response, customer energy data and vehicle charging. The platform's importer refused none and found 1,197 operations [9].

What that exercise taught us. The open energy interfaces are built for reading. By method, 754 of the 1,197 operations read and 443 are sent as changes. The British balancing market's data service alone carries 326 operations, and every one reads. The standards lean the other way: 18 of the demand response standard's 31 operations change something. One market platform sends 270 of its 301 operations as changes, though most of them are queries, so the method alone cannot tell a read from a write. Add nothing that writes to a control system at all.

What is not published. Control system, meter data, billing and trading suppliers publish no open description. Some asset management systems generate the description inside each installation, behind a sign-in; some market operators keep theirs behind registration. Where nothing is published, the company supplies its own description, or the agent reads from the data warehouse.

Operational networks. The safe design keeps the agent on the office side of the boundary with the operational network, reading from the historian's copy or a reporting database, never from the control system itself.

So choose the operations at import, start with the ones that read, and add nothing that changes a record until the approval step has been seen working.

PART 04

What an agent may do

One rule runs through every agent in this paper: it prepares, and a named person decides.

01	**It reads** what the person using it could read, and no more.
02	**It drafts** the work order, the visit report, the complaint response, the regulatory return, with each statement tied to its source.
03	**It lists** what is missing, what is out of limits, what falls outside the procedure.
04	**It stops** before anything that changes a record or reaches a customer, a crew, a regulator or a colleague. A tool that changes something waits for approval.
05	**It decides nothing about a person.** It does not decide a customer's debt, deposit or supply, does not score staff, and does not name anyone as the cause of an incident.

Four lines are drawn tighter than the rest.

Control systems.

No agent writes to a control system, operates a switch or dispatches plant or a crew. Grid and plant data reach it read only. An agent may check a switching schedule against the procedure and list every missing step, or propose which feeders to inspect first, with the readings behind it; the authorised engineer decides.

Trading.

An agent may prepare a briefing from published market data, with the source of each figure. It places no trade and sets no part of an order.

Customers and bills.

An agent may explain a bill, check a disputed bill against the meter reads and the tariff, and draft a response. It changes no bill, decides no redress, and disconnects no one. Text to a customer is sent by a person, and a customer can always reach a person.

Staff.

An agent may draft a training note from a procedure. It does not assess, rank or allocate people.

PART 05

What the rules already ask

This part needs a lawyer's read before it is relied on. It reports what the texts say; it is not advice.

Critical infrastructure

The trigger is a safety component. The European Union's AI Act lists as high-risk "AI systems intended to be used as safety components in the management and operation of critical digital infrastructure, road traffic, or in the supply of water, gas, heating or electricity" [10]. An agent that drafts a work order or answers a question on a procedure is not a safety component of the grid. One that could switch or dispatch plant might be. Which side a given agent falls on is a question for counsel.

The dates have moved. Regulation (EU) 2026/1744 set the high-risk duties to apply from "2 December 2027 as regards AI systems classified as high-risk pursuant to Article 6(2) and Annex III" and from "2 August 2028" for those tied to Annex I [11]. Critical infrastructure sits in Annex III, so its date is 2 December 2027.

A person must be able to oversee it. A high-risk system must be built so that it "can be effectively overseen by natural persons during the period in which they are in use" [10]. Critical infrastructure is spared the Act's fundamental rights impact assessment, "with the exception of high-risk AI systems intended to be used in the area listed in point 2 of Annex III" [10]. The oversight duty stays.

PART 05 · CONTINUED

Customers

People must be told. The AI Act asks that systems "intended to interact directly with natural persons" be built so that "the natural persons concerned are informed that they are interacting with an AI system", unless that is obvious [10]. In Great Britain, the energy regulator says its fair treatment rules apply "when these companies use AI", and that companies should "be clear and transparent when they are using AI" and "make it easy to speak to a human if needed" [12].

Scoring a customer can be high-risk. The AI Act lists systems that "evaluate the creditworthiness of natural persons or establish their credit score" as high-risk, and its recitals say such systems "determine those persons' access to financial resources or essential services such as housing, electricity, and telecommunication services" [10]. No agent in this paper scores a customer.

Decisions about a person stay with a person. Under the GDPR, a person "shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her", with stated exceptions [13].

The bill has required contents. The Electricity Directive's Annex I sets "MINIMUM REQUIREMENTS FOR BILLING AND BILLING INFORMATION", with key information that "shall be prominently displayed to final customers in their bills" [14]. An agent that explains a bill must not drop or change them.

Markets

An algorithm that sets orders is regulated. The EU's market integrity rules for wholesale energy, as amended in 2024, define algorithmic trading as trading "where a computer algorithm automatically determines individual parameters of orders to trade", "with limited human intervention or no such intervention at all" [15]. A firm that engages in it "shall have in place effective systems and risk controls" to keep its trading systems resilient and to "prevent the sending of erroneous orders to trade", and must notify its regulator [15]. An agent that prepares a market briefing sets no order. One that set an order's price or timing would fall to be read against this text.

Security and data

MARKET	WHAT WE READ	WHAT IT ASKS
European Union	NIS2 Directive	Lists electricity suppliers, distribution and transmission system operators and producers as sectors of high criticality; they must take "appropriate and proportionate technical, operational and organisational measures" for the security of their systems [16]
European Union	Network code on cybersecurity for electricity, Regulation (EU) 2024/1366	Sets "sector-specific rules for cybersecurity aspects of cross-border electricity flows", and can reach critical ICT service providers [17]
United Kingdom	Ofgem, "Approach to AI assurance in the energy sector"	A consultation published on 18 November 2026, open until 15 January 2027 [12]
Saudi Arabia	Operational technology cybersecurity controls (OTCC-1:2022)	Minimum cybersecurity requirements to protect industrial control systems, as "an extension to the NCA's Essential Cybersecurity Controls" [18]

We did not open the primary text of the North American reliability standards, Australia's critical infrastructure law, or the rules in the United Arab Emirates or India, and found no rule in any of them aimed at AI agents in energy. A company in those markets should ask its own counsel.

PART 05 · CONTINUED

What the rules have in common

01	**A person stays in charge** of anything that affects the safety of supply, a customer's account or a market order.
02	**Customers are told** when they deal with AI, and can reach a person.
03	**Control systems and trading are held to their own rules,** and an agent that reads and drafts sits outside them.
04	**The company knows where the data is** and who can reach the operational network.

PART 06

A worked example: one regional utility

Take a made-up regional utility that runs an electricity network and supplies water: about 2,000 people, some 600,000 customers, 40 depots. Nothing here is a customer's story, and every figure is illustrative; it is section 4 put to work.

01

In the field, a line worker asks what the operating procedure says about working near a damaged pole. The agent answers from the company's own procedure and shows the step. Asked about a situation the procedures do not cover, it says so and tells the worker to call the supervisor.

02

Before switching, an engineer uploads a planned switching schedule. The agent checks it against the switching procedure and lists one missing earth and a step out of order, each with the clause. The authorised engineer returns the schedule for correction.

03

At the permit office, a completed permit to work arrives for a pump station. The agent lists a missing isolation and an unsigned precaution against the site's permit procedure. The permit authority refuses the permit until both are fixed.

04

In asset management, the agent groups the quarter's inspection findings by asset and severity and lists 14 that are past their due date. The integrity engineer decides every action.

05

On the complaints desk, a customer disputes a winter bill. The agent sets the bill against the meter reads, the tariff and past bills, finds an estimated read, and drafts the response. The complaints handler approves it and decides any redress.

06

In regulation, the agent drafts a quarterly return from the reporting database and the regulator's template, with each figure traced to its source. The regulatory manager checks and submits it.

Every one of these reads documents or records the company already has, writes nothing to a control system, changes no record by itself, and leaves a record of who asked and what was shown.

PART 07

Seven questions before an agent touches the network

01 **Where does it run?**
On our infrastructure, in our private cloud, or in the supplier's service? Where is the model? Can it run with no connection to the internet?

02 **What can it read?**
Does it act as the person using it, with their permissions, or as one all-seeing account?

03 **What can it change?**
Which tools change a work order, an asset record, a customer's account or a return, and does each wait for a person?

04 **Can it reach a control system?**
Show us that no tool writes to the operational network, and which side of its boundary the agent sits on.

05 **Who decides?**
Show us the step where a named person approves, and what the agent does when nobody has.

06 **What is recorded?**
Who asked, which tool ran, what was decided, and for how long is that kept?

07 **Can we stop it?**
One agent, one tool, or all of them, at once, by us.

PART 08

Where to start

Pick work where reading is slow, the decision stays with a person, and no customer or control room waits on the agent.

Questions on your own procedures

Field workers ask what the company's own procedures and safety rules say and get the answer with the step cited. No operational system is touched, so it can run on the first day.

Permit to work checks

The agent lists every missing isolation, signature or precaution on a permit. The permit authority issues or refuses it.

Complaint responses

The agent sets out what happened and drafts the response. A complaints handler approves it.

One team, one region, invented data first. Then beside the manual process, measured against it.

About Katonic AI

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI/SOLUTIONS/ENERGY

SOURCES

Read on 22 November 2026. Figures are the publishers' own and were not audited by us.

01 Octopus Energy press release, "Octopus Energy's AI trial wins customer approval", 7 July 2026. octopus.energy

02 American Water Works Association, "'Rocket fuel' for employees: Inside a utility's AI journey", 18 September 2026. awwa.org

03 IT Brief Australia, "Shell expands C3 AI reliability deal across global assets", 10 June 2026. itbrief.com.au

04 Honeywell press release, "Honeywell Unveils Commercial Launch of AI-Powered Control Room Assistant Following Successful Pilot", 19 March 2026. honeywell.com

05 Honeywell press release on Experion Cognition, 9 June 2026. honeywell.com

06 Control Global, "Honeywell previews Experion Cognition release", 11 June 2026. controlglobal.com

07 IBM documentation, "Deploying watsonx.ai on-premises for AI Service" (Maximo Application Suite). ibm.com/docs

08 n8n workflow template 18491, "Detect smart meter anomalies and route critical alerts via HTTP API", n8n.io; Salesforce, Agentforce for energy and utilities customer service use case page, salesforce.com.

09 Template libraries of agent and workflow platforms, searched 22 November 2026; and 19 interface descriptions fetched from their publishers and standards groups the same day and run through the platform's importer (Katonic research).

10 Regulation (EU) 2024/1689 (AI Act), Annex III points 2 and 5(b); recitals on critical infrastructure and creditworthiness; Article 14(1); Article 27(1); Article 50(1).

11 Regulation (EU) 2026/1744, 8 July 2026, recital 40 and the amended Article 113.

12 Ofgem, "Understanding AI in the energy sector: the benefits, the risks and your rights", 5 February 2026; and "Approach to AI assurance in the energy sector", consultation published 18 November 2026. ofgem.gov.uk

13 Regulation (EU) 2016/679 (GDPR), Article 22(1).

14 Directive (EU) 2019/944 (Electricity Directive), Annex I point 1.1.

15 Regulation (EU) No 1227/2011 (REMIT) as amended by Regulation (EU) 2024/1106, Article 2 definition of algorithmic trading and Article 5a(1) and (2).

16 Directive (EU) 2022/2555 (NIS2), Article 21(1) and Annex I point 1(a).

17 Regulation (EU) 2024/1366 (network code on cybersecurity aspects of cross-border electricity flows), Articles 1 and 2.

18 National Cybersecurity Authority (Saudi Arabia), Operational Technology Cybersecurity Controls (OTCC-1:2022) page. nca.gov.sa