

Capacity is not the only energy bottleneck

Proving connected energy services can work, scale and recover when it matters



A viewpoint for functional leaders across engineering, quality, product, digital, DevOps and OT working on connected energy services.

The electricity revolution is not simply a capacity story. It is also creating a much more connected, software-dependent operating system. Rising electricity demand from AI and data centres, electric transport and heating, and industrial growth is accelerating investment in generation, networks, storage and flexibility. Yet functional leaders face a second constraint: proving that the connected services coordinating the system can work, change, scale, recover and remain trusted as one.

3.6%

Expected average annual growth in global electricity demand between 2026 and 2030 – around 50% faster than the previous decade. ¹

2x

Electricity demand from data centres worldwide is expected to more than double by 2030, with AI the most significant driver. ²

The energy sector's public conversation is dominated by capacity. How quickly can new generation be built? How fast can networks be reinforced? Where will data centres, storage and other large loads connect? Those questions are critical, but they are not the whole engineering challenge.

The scale of demand growth explains the urgency. The International Energy Agency expects global electricity demand to increase by an average of 3.6% a year between 2026 and 2030 – around 50% faster than the average rate of the previous decade. It also expects electricity demand from data centres worldwide to more than double by 2030, with AI the most significant driver of that increase.

Demand growth does not have one cause. The electrification of heat and transport describes the specific shift from oil and gas towards electricity at the point of use. AI and data centres are a separate and increasingly significant driver. Industry, cooling and new digital services add further pressure. At the same time, more weather-dependent generation, storage, distributed assets and flexible demand increase the need for forecasting, orchestration and rapid operational response.

These changes do not simply add megawatts. They add software, data, devices, interfaces, suppliers and operational decisions. The electricity system is becoming larger, but it is also becoming more distributed, more responsive and more dependent on connected services that must coordinate physical and digital behaviour in real time.


RESILLION VIEWPOINT

The next bottleneck may not be capacity itself. It may be confidence in the connected services that turn capacity into a reliable customer, operational or market outcome.

The demand story is only half the story

A higher-demand electricity system needs more infrastructure. It also needs confidence in the services that control, balance and monetise that infrastructure. A smart tariff, flexibility dispatch, battery-control instruction or vehicle-to-grid transaction can cross customer applications, cloud platforms, identity services, communications networks, market systems, operational technology and physical assets before the intended outcome is achieved.

ONE OUTCOME, SEVEN HAND-OFFS



The intended outcome is only achieved once every hand-off has completed with the right timing and state.

ILLUSTRATIVE SCENARIO


Tens of thousands of charge points, one demand event

Consider an illustrative demand event involving tens of thousands of connected charge points. The service provider's platform remains available. Its APIs meet their response targets. Each delivery team reports that its own component is operating normally. Yet a supplier interface begins to throttle requests, retry traffic increases the load and some devices respond too late. When the platform recovers, a small number of duplicate commands appear.

Nothing has failed cleanly. Most individual dashboards may still look healthy. But the connected service is no longer delivering a trustworthy outcome. The issue sits in timing, state, capacity and recovery across the hand-offs –precisely the areas that component-level assurance can miss.

A technical pass can coexist with an operational failure.

That distinction matters because the unit of delivery is no longer one application, device or operational system. It is a service outcome that crosses technical and organisational boundaries. The assurance boundary therefore has to follow the customer, operational or market result –not the organisation chart, delivery workstream or supplier contract.



The real risk often sits in the hand-offs

Connected energy services rarely fail because one component stops working completely. More often, the weakness appears in the hand-offs between systems, suppliers and operational processes.

- A cloud platform may be available while the operational application cannot authenticate.
- A device may conform to the expected standard but respond too slowly under real network conditions.
- A supplier may restore its service while transactions, telemetry or device state remain inconsistent.
- A security control may successfully contain a suspected compromise but also remove the access route needed for recovery.

These are interface failures rather than isolated component failures, and they are where fragmented assurance is most likely to miss the real operational risk.

Traditional assurance is often organised around projects, components and delivery teams. Standards, certifications and supplier evidence provide essential foundations, but they cannot on their own prove that the deployed versions, configurations, network conditions and operating assumptions will work together.

FOUR WAYS THE SAME OUTCOME CAN FAIL

Interoperability

Interoperability exposes the distinction. A device may conform to a protocol and still behave differently when messages are delayed, duplicated or received out of order. A cloud platform may scale in isolation while a gateway, communications network or device estate becomes the true constraint. The component has passed; the outcome has not been proven.

Performance and recovery

Performance and recovery also have to be treated as end-to-end properties. An API can meet its target while the customer journey remains too slow. A platform can accept instructions faster than devices can process them. Failover can restore infrastructure while leaving transactions, telemetry or device state inconsistent. A restarted platform is not necessarily a recovered service.

Software change

Software change creates the same problem. A small code or configuration update can alter billing, settlement, customer communications, operational control or third-party integrations. Test completion is a weak proxy for readiness when the evidence does not show how the complete service behaves in the conditions that will actually be operated.

Cyber resilience

Cyber resilience belongs in the same engineering conversation. Prevention matters, but so do detection, containment, essential operation and the re-establishment of trust. After an incident, leaders need evidence that identities, configuration, data, firmware, device state and service behaviour are trustworthy again – not simply that systems are available.

These are not four unrelated assurance disciplines. They are different ways in which the same connected outcome can fail. Treating them separately can create a reassuring collection of component-level passes while leaving the most important question unanswered: will the deployed combination work as one service?



Proof in practice: Octopus Energy and whole-service assurance

Practical energy programmes show why this whole-service view matters. Through the UK Government's Interoperable Demand Side Response programme, a Resillion-led consortium brought together appliances, service platforms and multiple organisations to demonstrate interoperability in settings indicative of the real world not simply conformance within each product. (3)

The same principle is visible in Resillion's work with Octopus Energy on its Vehicle-to-Grid. Resillion helped validate the connected combination of hardware, software, interoperability and regulatory requirements before deployment. The assurance work considered the vehicle, charger, grid-connection equipment and control software as part of one connected proposition, helping reduce deployment and integration risk and support the evidence needed for certification. (4)

The point is not simply that each technology had been tested individually. Confidence depended on whether the components could operate together as the intended service under realistic conditions exactly the distinction between component-level assurance and Engineering Confidence.

The wider lesson is uncomfortable but useful: more assurance activity does not automatically create more confidence. Confidence improves when evidence is current, representative, traceable to the service outcome and explicit about what remains uncertain.

“Every component may have passed, but have the hand-offs between those components been exercised under the conditions that matter including peak demand, degraded communications, supplier delay, software change, containment and recovery? Engineering Confidence depends on proving the connections between systems, not only the systems themselves.”

Bill Chard, Head of Energy Sector, Resillion

Engineering confidence must follow the outcome

Resillion uses the term engineering confidence to describe the justified belief that a connected service will achieve its intended outcome within agreed tolerances under normal, peak, degraded and recovery conditions. It's not a promise of certainty, and it is not another governance layer. It's a clearer way to connect service risk, evidence and engineering decisions.

Four questions functional leaders should ask

1

Can the ecosystem work as one service?
Start with the complete customer, operational or market outcome and map the systems, suppliers, interfaces, devices, data and manual processes that determine it. Focus particularly on timing, state, versioning, error handling and ownership at the boundaries.

2

Can change be introduced at the required pace?
Automation should protect critical journeys, data flows, configurations and known failure paths, while specialist judgement concentrates on new integrations, physical assets and unusual operating conditions. Readiness is reached when the evidence supports the intended outcome within a defined boundary - not simply when every delivery activity is complete.

3

Can the service perform, degrade and recover?
Leaders need to know where the complete service saturates first, what minimum outcome is acceptable during disruption, how retries and backlogs behave, and how transactions, telemetry and device state will be reconciled. Recovery is not complete until the service has returned to a trustworthy state.

4

Can the service remain secure and trusted?
The organisation must be able to detect abnormal behaviour, contain affected identities, devices or suppliers, preserve essential operation and re-establish trust across the full service. Availability alone is not sufficient evidence of recovery.

But the questions alone are not enough. Each one should be reviewed to achieve an evidence-backed operating decision, not another set of separate technical sign-offs.

Turning answers into engineering confidence



Outcome
What complete customer, operational or market outcome must remain within tolerance?



Dependencies
Which systems, suppliers, devices, data and interfaces determine that outcome?



Conditions
Which normal, peak, degraded, recovery, change and attack conditions could undermine it?



Evidence
What has actually been demonstrated in the versions and configurations that will operate?



Decision
What decision does that evidence support: proceed, proceed with constraints, delay or stop?

Evidence matters only when it changes the decision

Engineering teams already generate plans, test results, dashboards, security findings and supplier reports. The difficulty is rarely a total absence of information. It is that the evidence is fragmented, differently scoped and often interpreted only within the team that produced it.

A useful whole-service review separates what has been demonstrated from what is believed, assumed or still unknown. It then connects that position to one explicit recommendation. The evidence may support proceeding as planned. It may justify proceeding with constraints for one supplier cohort, device range,

volume or operating window. It may require delay while a material gap is closed, or stopping because the service state cannot be trusted or the exposure cannot be controlled.

Return to the charge-point scenario. If recovery testing shows that duplicate commands and slow state reconciliation affect one supplier combination, a defensible decision is not a blanket approval or an indefinite stop. The proven cohort can continue within a controlled boundary and enhanced monitoring, while the affected combination remains excluded until its recovery behaviour meets the agreed tolerance.

ONE EXPLICIT RECOMMENDATION

Proceed

as planned

Proceed with constraints

for one supplier cohort, device range, volume or operating window

Delay

while a material gap is closed

Stop

because the service state cannot be trusted or the exposure cannot be controlled

A constrained decision is not a failure of assurance. It is often the strongest decision available because it reflects the exact boundary of what has been proven.

Energy organisations do not need an enterprise-wide transformation programme before they can improve confidence. They can start with one high-consequence connected service. Define the outcome. Map the dependencies. Exercise normal, peak, degraded and recovery conditions. Separate representative evidence from assumption. Agree the safe release or operating boundary, and use live outcomes to strengthen the next decision.

The energy system will need more generation, network capacity, storage and flexibility. It will also need a more disciplined view of what has actually been proven. The next bottleneck will not always be a physical constraint. It may be the gap between the assurance activity teams have completed and the decision leaders can responsibly make.

Organisations that follow the outcome, exercise the deployed combination and connect evidence to a clear decision will be better placed to move at pace without mistaking activity for confidence.

Start with one critical service

A Resillion workshop helps cross-functional teams define the outcome, expose dependencies, test high-consequence scenarios and establish the evidence required for the next decision.

01

Define the outcome

02

Expose dependencies

03

Test high consequence scenarios

04

Establish the evidence for the next decision

SOURCES

1. International Energy Agency, Electricity 2026 -Executive summary -global electricity-demand outlook for 2026-2030
2. International Energy Agency, Energy and AI -data-centre and AI electricity-demand outlook to 2030
3. UK Government, Interoperable Demand Side Response programme -programme structure and real-world demonstration
4. Resillion, Octopus Energy Vehicle-to-Grid case study -vehicle-to-grid certification and end-to-end assurance

Editorial note: Engineering confidence is a Resillion working point of view, not a certification scheme or formal industry standard. Market figures were checked in August 2026.