Software Improvement Group

Energy signals 2025

The software risks and opportunities shaping the future of energy

Table of contents

Executive summary

The 7 software signals shaping energy in 2025

As energy providers race to decarbonize, modernize, and digitize, software quality is becoming a critical but often overlooked factor in achieving these goals.

This report draws on Software Improvement Group (SIG)'s analysis of over 300 billion lines of code to surface the most urgent risks and opportunities facing the industry. From security shortfalls to AI adoption challenges, the findings highlight how software can accelerate or undermine transformation.

1. Security gaps threaten resilience

→ 67% of energy systems fall below the average benchmark for software security, exposing them to operational disruption, compliance risk, and cyberattack.

2. Poor build quality increases operating costs

→ 42% of energy systems fall below SIG's recommended build quality threshold–worse than the cross-industry average. Poor build quality reduces innovation capacity and drives up annual maintenance effort and cost.

3. Green IT is now essential infrastructure

→ Energy companies must address software's hidden footprint. Optimizing code structure and refactoring can reduce energy use by up to 90%.

4. AI is powerful but fragile

→ Among large companies, AI adoption rates rose from slightly over 15% of firms using AI in 2020 to nearly 40% in 2024. However, 73% of AI systems show structural quality issues. Without robust software foundations, AI becomes hard to scale and expensive to maintain.

5. Legacy systems undermine change

→ Systems with poor architecture are 40% slower to evolve. Despite an average architecture rating of 4 stars in energy, legacy platforms still pose cost, security, and scalability risks.

6. Cloud readiness is a competitive advantage

Energy companies outperform other industries on cloud readiness (3.9 stars vs. 3.1)*, but technical debt still holds some systems back. True readiness requires modular design, automation, and distributed knowledge.

7. Knowledge is a hidden risk

59% of energy companies meet SIG's recommended knowledge distribution threshold—twice the cross-industry norm. But without continuous documentation and mentoring, knowledge monopolies can silently increase costs, delays, and outages.

Software quality is no longer optional

It is the foundation for securing infrastructure, reducing emissions, scaling innovation, and delivering affordable, reliable energy. Energy companies that invest in code quality, architecture, and team resilience will be best positioned to lead the next era of transformation.

**A note on SIG's star ratings:*

SIG's star ratings are a way to evaluate and compare software quality.

The ratings range from 1 to 5 stars, with half-star increments (so technically 0.5 to 5.5). A 3-star rating represents the market average. The ratings are based on comparing a system's properties to SIG's benchmark, which includes thousands of systems and is recalibrated yearly.

Foreword

The energy industry stands at the heart of one of the most complex transformations of our time. As the global push for decarbonization intensifies, energy providers are under immense pressure to modernize infrastructure, integrate renewables, and meet rising demand—all while keeping the lights on.

At the core of this challenge lies software. Digital systems have long been critical to how energy is produced, managed, and delivered. But as digital complexity increases, so do the risks. Cyberattacks on critical infrastructure are surging. Legacy software is driving up costs. And the environmental impact of IT is under growing scrutiny.

For energy providers, success depends on innovation and efficiency. In a market where margins are tight and competition is fierce, keeping costs down is non-negotiable. That's why software quality matters. It's what enables organizations to:

- Protect critical infrastructure from growing cyber threats,
- Avoid costly outages through stronger software build quality,
- Reduce emissions and energy use through greener code,
- Deploy AI responsibly while controlling complexity,
- Modernize legacy systems without spiraling costs or delays.

This report draws on our proprietary database—covering over 300 billion lines of code across 20,000 systems—to deliver data-driven insights on the biggest software challenges facing the energy sector today.

If you're striving to compete, lower costs, and build the systems that will power a cleaner future, this report will help you make software a source of strength, not risk.

Luc Brandts

CEO
Software Improvement Group

Chapter 1: Cybersecurity shortfalls threaten energy resilience

Key findings:

- **67% of systems in the energy sector have a below-average degree of security controls.**
- The **average system has 19 security findings**, increasing exposure to operational disruptions, regulatory breaches, and cyberattacks.
- **50% of enterprise software systems show vulnerabilities** in open-source components each month, highlighting the ongoing risk from unmanaged third-party dependencies.

Rising cyber threats target the energy sector

The energy sector has become a top target for cybercriminals and nation-state actors alike. With the convergence of IT and operational technology (OT) across smart grids, connected substations, and cross-border systems, the industry faces growing exposure to both digital and physical disruption.

According to the World Economic Forum (2025), the global energy sector is among the top five industries at risk of nation-state cyberattacks over the next 24 months. Microsoft reports a 40% rise in attacks on critical infrastructure worldwide, and within Europe, cyberattacks on the power sector increased by 57% between 2021 and 2023—primarily aimed at grid operators and cross-border control systems. (KnowBe4, 2025)

The operational impact is already being felt:

45% of European energy companies have experienced at least one cyberattack that affected operations, including shutdowns, delays, or data breaches (KnowBe4, 2025).

Cyberattacks on energy infrastructure don't just disrupt business—they threaten national security and public safety. As energy systems increasingly combine physical and digital components, breaches can cause cascading failures. A successful coordinated attack could result in multi-day regional blackouts, damage to connected home devices, and disruption to critical services such as heating, water, and telecoms.

Beyond the operational risks, the **financial consequences are severe.** The average cost of a successful cyberattack on a European utility is estimated at **€7.35 million**—and that excludes regulatory fines, legal fees, and reputational damage (Eurelectric, 2025). For energy providers under pressure to keep prices low and remain competitive, every breach drives up operating costs and erodes the ability to deliver affordable, reliable service to consumers.

Notable incidents from the past highlight the stakes:

- The Colonial Pipeline attack (2021) shut down the largest fuel pipeline in the U.S., disrupting supply to 17 states and resulting in $4.4 million in ransom payments.
- A 2015 attack on Ukraine's power grid caused widespread outages for hundreds of thousands of citizens, demonstrating the real-world consequences of digital vulnerabilities.

How are security risks introduced?

These incidents raise a critical question: where do these vulnerabilities begin?

It's tempting to think of cyberattacks as the work of elite hackers breaching systems with cutting-edge tools. But in reality, most attackers are simply looking for an open door, and those doors are often hidden in the code.

Security risks typically originate during development: through unpatched dependencies, poor coding practices, or weak architecture. If software isn't built with security in mind from the start, these hidden flaws become attack surfaces that continuously expose businesses to disruption, data loss, and regulatory risk.

Below-average security ratings are leaving the energy sector exposed

Data shows that **67% of energy systems have a security rating below 3.5 stars**, indicating a below-average level of software security controls.

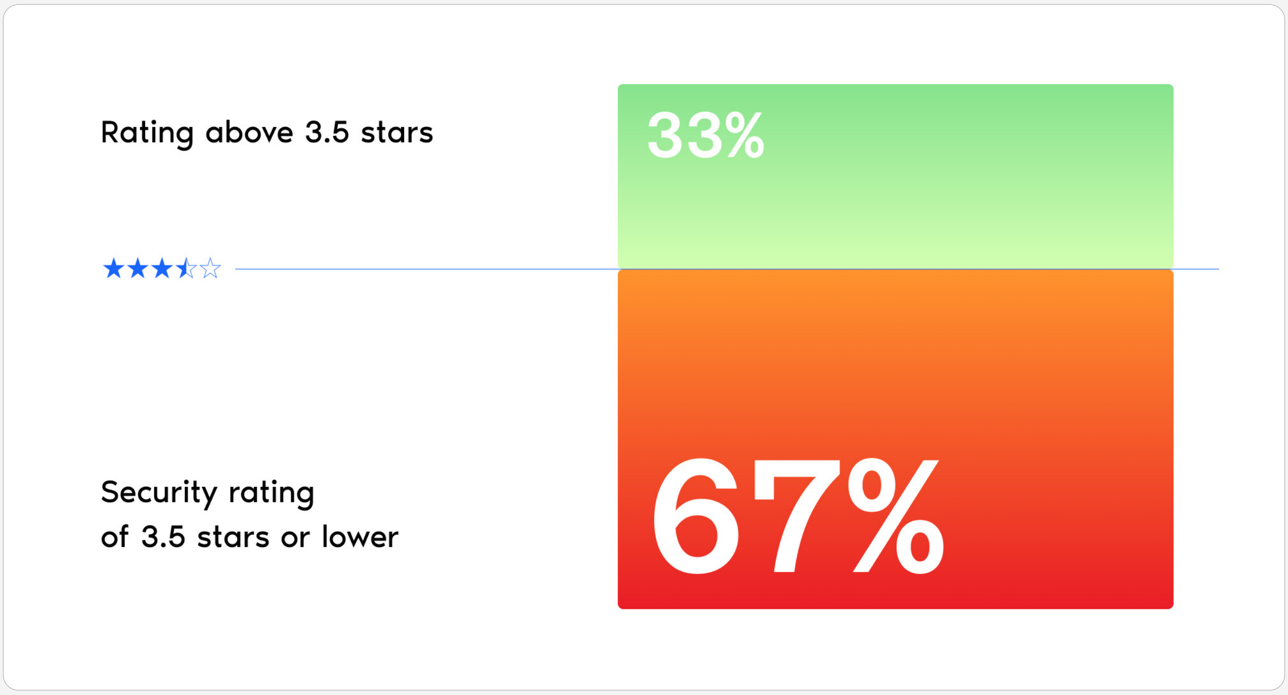

Based on a snapshot of active security findings in all systems in our data warehouse on a random day medio 2023.
Our SAST (Static Application Security Testing) security model that ranks software systems from 1 to 5 stars.
We evaluate system properties through a thorough analysis of the source code and infrastructure. This includes reviewing the codebase and other artifacts (such as documentation) to derive scores for various system characteristics. These characteristic scores are then mapped to the OWASP Top 10, which identifies the ten most critical risks in web application security. This report is compiled by a group of security professionals from around the globe and serves as an "awareness document." It is recommended that all organizations integrate its findings into their processes to effectively reduce and manage security risks.
The star rating reflects your compliance benchmark against the OWASP Top 10: 1. Severely low degree of security controls, 2. Very low degree of security controls, 3. Low degree of security controls, 4. Moderate degree of security controls. 5. High degree of security controls.
It is important to note that a 4- or 5-star rating does not guarantee full security; it simply indicates that security considerations have been factored into the design and implementation, making vulnerabilities less likely.

On average, each system in our database contains **19 security findings**, many of which can expose systems to operational and regulatory risk.

19

Critical security findings per system

*This estimation is based on a snapshot of active security findings in all systems on a random day medio 2023. The number of findings were then translated into an average of security findings (1.16) per person year (size of system), which was then used to calculate an estimation of critical security findings per system.
A software system refers to a collection of interrelated programs, data, and documentation that work together to perform specific tasks or functions and have their own team. For example, a single application can consist of multiple interconnected systems. The size of the system we took as an average equals 16.3 person years which indicates how many years it would take a single person to rebuild the same system from scratch.

This number reflects an average per system, based on a typical-sized system in our benchmark. However, it's important to note that financial services institution (FSI) systems can be up to ten times larger than the average system in our benchmark. Generally, larger systems tend to have lower security ratings, which correspond to a relatively higher number of security findings, while smaller systems often achieve higher security ratings.

We evaluate system properties through a thorough analysis of the source code and infrastructure. This includes reviewing the codebase and other artifacts (such as documentation) to derive scores for various system characteristics. These characteristic scores are then mapped to the OWASP Top 10, which identifies the ten most critical risks in web application security. This report is compiled by a group of security professionals from around the globe and serves as an "awareness document." It is recommended that all organizations integrate its findings into their processes to effectively reduce and manage security risks.

Not every security flaw turns into a breach, but with, the cost of an attack on a european utility estimated at €7.35 million and the average breach globally costing $4.88 million, why take the risk? Catching vulnerabilities early in the development process can help organizations avoid things like costly breaches, business disruption, and reputational harm.

With attack surfaces expanding and critical infrastructure at risk, companies can't afford to treat software security as an afterthought. Identifying and remediating vulnerabilities should be a priority for every energy provider aiming to maintain operational resilience and regulatory compliance.

The open-source dilemma

Open-source software (OSS) usage is on the rise, with 96% of organizations globally increasing or maintaining their use of OSS (OpenLogic, 2025). Certainly, OSS helps cut costs, develop and deploy new applications faster, and modify and tailor software, but it also introduces hidden risks.

Our earlier findings showed that 50-60% of enterprise software systems contain at least one vulnerable open-source dependency each month, and 30% have a critically vulnerable dependency.

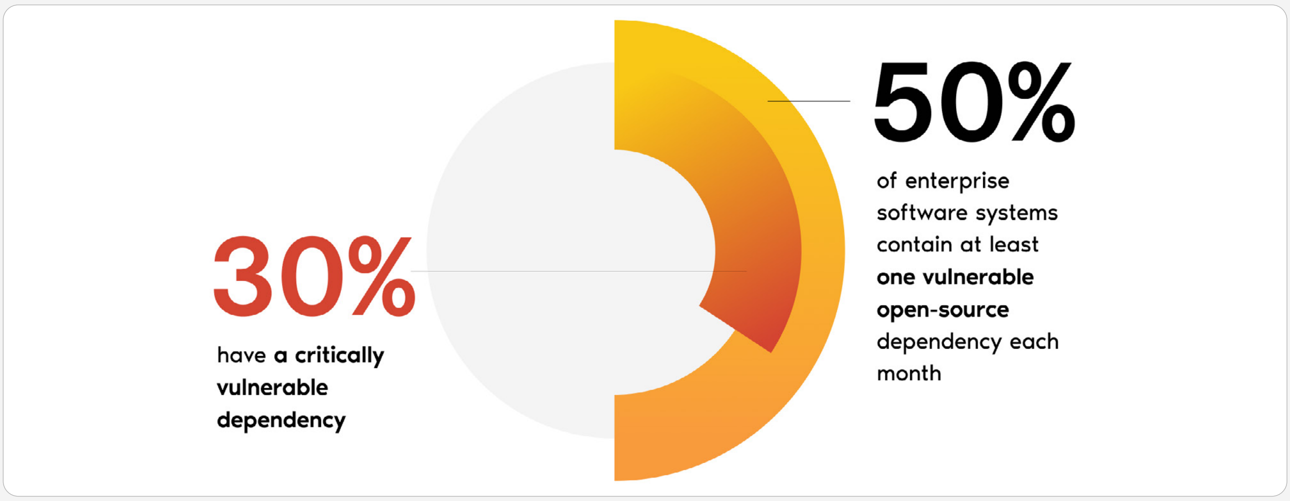

The multi-layered approach to cybersecurity

To establish a strong cybersecurity posture, organizations need a layered approach that combines multiple security measures.

Three key methodologies in software security testing include:

- **Penetration Testing (Pentest) Simulates external attacks to uncover vulnerabilities.**
- **Static Application Security Testing (SAST) Analyzes the source code to detect weaknesses before deployment.**
- **Software Composition Analysis (SCA) Scans third-party open-source libraries and dependencies for known vulnerabilities.**

No single method is enough on its own. Both SAST and SCA are complemented by penetration testing to form a complete security assessment. Together, these techniques enable earlier detection, stronger compliance, and more secure software from the start.

Energy companies must take a proactive stance. That means implementing Software Composition Analysis (SCA), patching known issues quickly, and treating OSS governance as a core part of their cybersecurity strategy.

Chapter 2: Poor build quality disrupts innovation

Key findings:

- **42% of energy systems fall below SIG's recommended build quality threshold,** compared to 37.4% across other industries.
- The **average system in energy requires just 3.9 FTE to maintain,** far less than the 13.4 FTE typical in other sectors–reflecting a higher share of smaller systems.
- **4-star systems unlock 30% more innovation capacity,** while **2-star systems face a 40% capacity shortage,** as teams are constantly firefighting instead of innovating.

Build quality issues put reliability at risk

Security isn't the only factor that can lead to outages. As smart grid environments grow more complex, the quality of the underlying software matters more than ever. Studies have identified over 30 different types of software or integration faults–from data syncing failures to interface mismatches–that can disrupt operations, compromise safety, or cause expensive downtime.

That's just one of the reasons a consistent overview of build quality is essential in the energy sector. Over the past decade energy companies have adapted fast to emerging technologies, driving operating costs down by as much as **60%** (McKinsey, 2022). That same cost discipline must now extend to the software layer, where poor build quality leads to higher maintenance costs, more outages, and slower delivery of innovation.

Energy lags behind on build quality

SIG analysis shows that **42% of systems in the energy sector fall below our recommended build quality**–a higher rate than the cross-industry average of 37.4%. These systems are harder to maintain, more prone to bugs, and more expensive to operate.

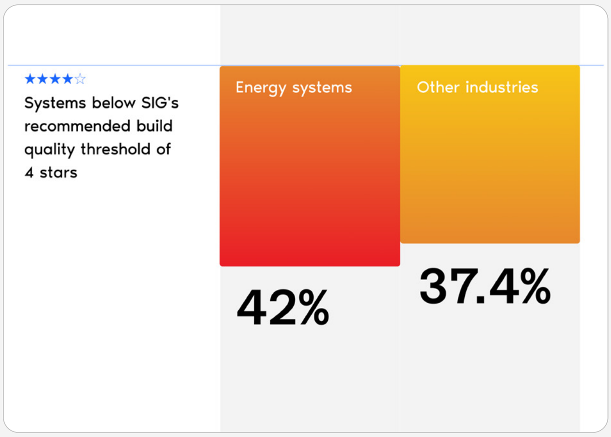

The result? Slower delivery of new features, and fewer resources available for strategic innovation.

The energy sector is uniquely positioned to act

There is good news. Compared to other sectors, energy systems in our database are smaller and more modular. The **median system in energy requires just 3.9 FTEs of maintenance effort**, compared to 13.4 in other industries. This could give energy companies a head start. In theory, improving even a small number of systems can have an outsized impact.

Average system maintenance effort

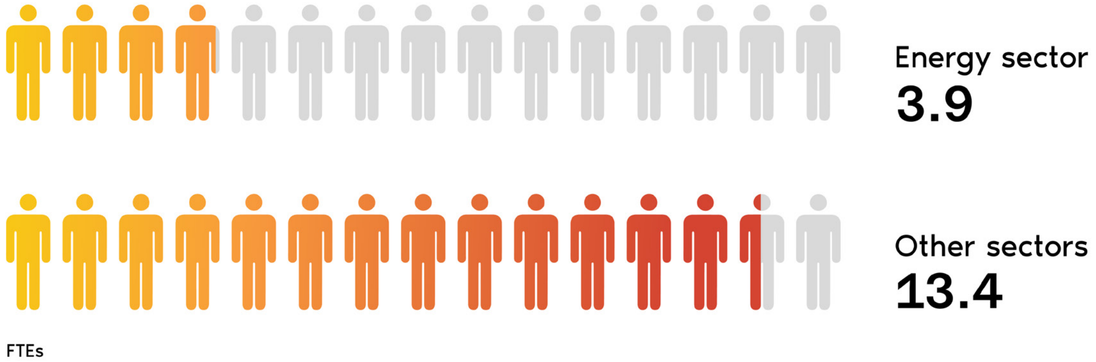

How do we measure system maintenance effort?
An FTE (Full-Time Equivalent) represents the effort of one full-time developer over a year. SIG uses this metric to express the maintenance capacity needed to keep a system at its current quality level. For example, 2 FTE means two full-time developers are required annually.

Better code unlocks capacity and cuts costs

Build quality isn't just about stability. It directly impacts a company's ability to evolve and innovate.

- **4-star systems** offer 30% more capacity for innovation and improvement.
- **2-star systems** experience a 40% shortage, as teams are forced to firefight instead of moving forward.

The cost impact is real. Our most recent State of Software Report estimates that poor build quality can increase system-level maintenance costs by up to €250,000–a hidden tax on every new feature or integration.

★★☆☆☆
CODE QUALITY
40% capacity shortage
for regular maintenance
★★★☆☆
CODE QUALITY
Market-average maintainability
★★★★☆
CODE QUALITY
30% extra capacity
for innovation and improvement

Earlier research has proven that systems with a 4-star maintainability score gain 30% extra capacity for innovation and improvement compared to 3-star systems. While 2-star systems lead to a 40% capacity shortage due to regular maintenance.

Cost optimization through software quality

Energy providers face constant pressure to reduce operating costs while maintaining grid reliability and investing in renewables. Software plays a central role in this balancing act, but poor build quality creates hidden costs.

- **Hard-to-maintain systems increase firefighting and reduce delivery speed**
- **Outages, defects, and reactive fixes raise operational overhead**
- **Cyber incidents cost millions per breach–not including recovery or penalties**

Improving build quality is a practical way to unlock capacity, reduce maintenance effort, and lower total IT costs without compromising performance or safety. With smaller, more modular systems than other industries, the energy sector is well-positioned to turn software quality into a competitive advantage.

Chapter 3: Green IT is non-negotiable in energy

Key findings:

- **Energy companies depend on IT systems that consume significant energy** and contribute to emissions.
- **Python,** despite being one of the most energy-intensive languages, is the most used in the energy sector.
- **Refactoring inefficient code can reduce energy usage by 17–90%,** depending on the system.

IT isn't just a back-office function– it's part of the energy footprint

As digital demand grows, the energy impact of software itself is becoming harder to ignore. According to Goldman Sachs, data center power consumption could rise by 160% by 2030, driven largely by AI workloads. Meanwhile, 40% of energy-related CO_2 emissions still come from fossil fuels used in electricity generation (World Nuclear Association, 2024). That makes every wasted watt count.

While the energy sector has made major strides in renewables and efficiency, it can't ignore what's happening in its software stack. Inefficient code, legacy systems, and bloated architectures silently drive up consumption and cost.

Sustainable software is a cost and compliance issue

Green IT isn't just about optics or Environmental, Social, and Governance (ESG) reporting. Optimizing software for energy efficiency has a direct impact on cost and compliance. As more energy companies commit to net-zero or emissions reduction targets, software becomes part of the measurable equation.

Sustainable software delivers measurable business value:

- **Reduced emissions** → Cut unnecessary compute, storage, and load to shrink IT's carbon footprint.
- **Faster software** → Lean systems improve performance and reduce latency.
- **Scalable by design** → Right-sized code adapts to real demand, preventing overprovisioning.
- **Ready for regulation** → Sustainable software supports ESG, CSRD, and IFRS compliance.

Programming languages impact energy use

The choice of programming language directly affects a system's energy consumption. Energy companies should consider sustainability alongside performance, scalability, and long-term build quality when selecting technologies. While rebuilding a system in a new language isn't always feasible, the trade-offs are worth evaluating–especially as digital workloads continue to grow.

Programming languages energy use

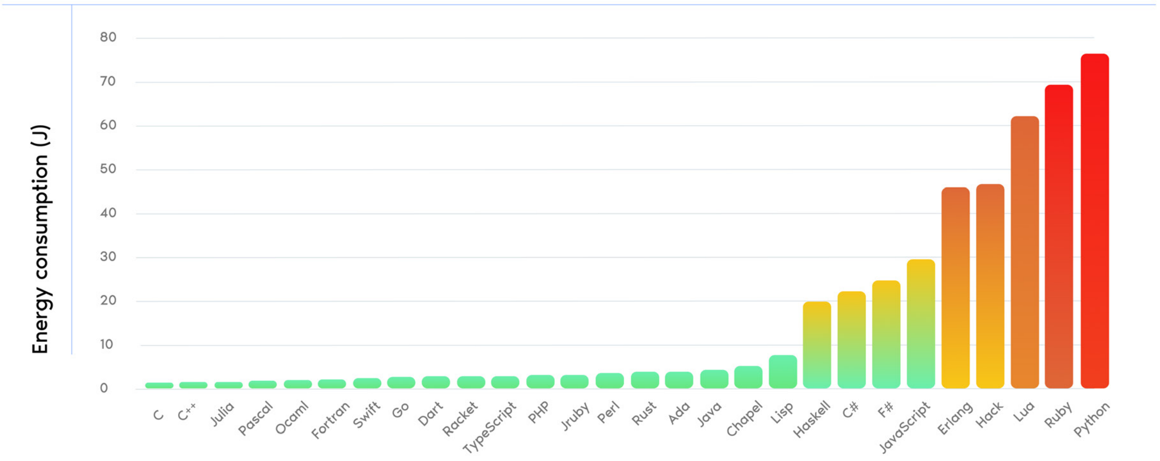

The energy sector must act on software efficiency

Python is the leading language in the energy sector, making up 18% of all systems in our database. Its popularity is understandable: Python consistently ranks among the top five most-used programming languages worldwide, thanks to its versatility and accessibility. For energy companies, it's especially well suited to AI, data analytics, and automation–core focus areas for the industry.

But Python's strengths come with trade-offs. It's also one of the most energy-intensive programming languages in use today.

Top 5

most used technologies

1. Java
2. C#
3. Typescript
4. Javascript
5. Python

This makes build quality even more important. Poorly structured Python systems lead to higher compute usage, costs, and emissions. But this is not inevitable. With disciplined engineering, Python can support scalable, maintainable systems. Performance can also be improved through techniques like ahead-of-time (AOT) compilation, which has been shown to make Python up to 19 times more energy efficient.

Optimizing inefficient Python code through better structure, algorithms, and architectural decisions offers a practical path to reducing resource consumption without changing languages or compromising flexibility.

Optimizing energy software for sustainability

At Software Improvement Group (SIG), we've been researching IT sustainability since 2014. Our findings consistently show a clear link between software quality and energy efficiency.

Simple code refactoring can slash energy uset

Our analysis shows that even without a language switch, **simple refactoring** of inefficient code can have a significant impact. On average, refactoring leads to a **17% reduction in energy consumption**, while in some extreme cases, **algorithmic improvements have slashed energy usage by up to 90%.**

Code refactoring energy use

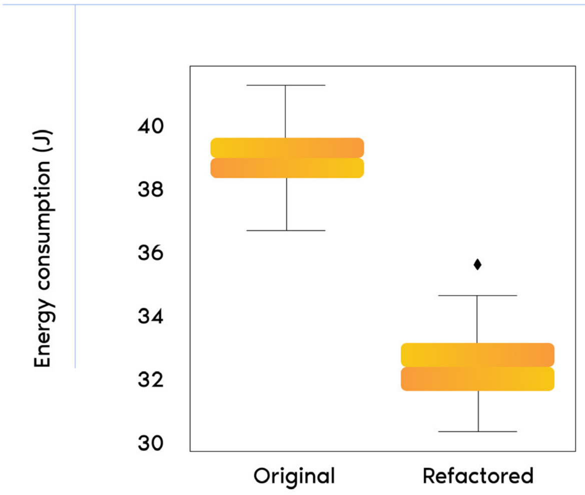

Refactoring performed: Replacing old standard Java collections with new counterparts

Average reduction of **17% in energy consumption**

As those in the field know best, energy transition isn't just about switching to renewables. It's also about making smarter use of the energy we already consume. That includes the unseen footprint of software. Optimizing IT systems for energy efficiency cuts emissions, reduces operating costs, strengthens compliance and improves performance.

In a sector under pressure to lead on climate and deliver affordable energy, Green IT isn't just responsible–it's essential.

Chapter 4: The hidden risks behind AI adoption in energy

Key findings:

- **73% of AI and big data systems have structural quality issues**, such as low maintainability, weak architecture, or missing controls.
- **85%** of U.S. energy leaders say their company is actively using or piloting AI technologies (Honeywell, 2025).
- As AI use grows, **software quality will be the bottleneck, not the model itself.**

AI is central to the future of energy

AI is already transforming the energy sector. According to Honeywell, **85% of U.S. energy leaders say their companies are actively using or piloting AI technologies**. The global market for AI in energy is projected to reach nearly $60 billion by 2030, with the biggest value areas in cybersecurity, predictive maintenance, and operational optimization.

But this momentum comes with serious questions about the long-term impact. To illustrate, training and running large AI models is energy intensive, and data center power consumption is projected to rise sharply, with some estimates suggesting **AI-related demand could account for up to 21% of global electricity use by 2030** (MIT, 2023).

This environmental cost is only part of the picture. AI can become a key enabler of sustainability. Research by Google and BCG suggests AI could help mitigate 5–10% of global greenhouse gas emissions by 2030 through smarter infrastructure and energy usage.

AI can become a key enabler of sustainability: From optimizing data center energy use and cloud workloads to supporting predictive maintenance and improving code efficiency, AI offers powerful ways to reduce IT's carbon footprint and extend hardware lifespans. In this way, AI can help conserve energy, rather than just consuming it.

Scaling AI requires more than good models

As AI adoption expands, so do its energy and resource demands. Global data center electricity use linked to AI could exceed **1,000 terawatt-hours (TWh)** annually by the end of the decade, roughly equivalent to the power consumption of Japan (WEF, 2024).

The systems driving this demand in the energy sector encompass far more than just generative AI. From predicting equipment failure and managing load balancing, to automating operations and analyzing smart grid data, energy companies are using AI to improve efficiency and decision-making at scale. With so many mission-critical applications, the complexity of managing and scaling AI systems continues to grow.

Yet building working AI models isn't enough to meet these ambitions. The real challenge is ensuring that the AI systems surrounding those models are efficient, stable, and scalable. In the energy sector, this means leveraging AI to drive smarter, more sustainable operations, whilst also making sure that the AI itself is engineered for long-term performance and minimal environmental impact.

Here, software quality is the difference between AI success and spiraling complexity. Poorly designed or bloated systems risk not only failure in production but also unnecessary environmental impact. To truly scale AI in a sustainable way and meet their ESG goals, energy companies must apply the same engineering discipline to AI that they apply to any other core system.

AI vs. traditional software

To address these new challenges, it's important to understand how AI systems differ from traditional software. Unlike conventional software, which follows fixed, pre-programmed rules, AI learns, evolves, and makes autonomous decisions. According to ISO/IEC 5338, co-developed by SIG, AI is classified as a software system with unique characteristics. These include the ability to think autonomously, learn from data, make decisions based on that data, and even potentially talk, see, listen, and move.

AI also needs regular retraining to stay accurate. What sets AI software apart from traditional software is that it doesn't just follow fixed rules. Instead, it learns by analyzing large sets of data, finding patterns, and using that knowledge to make educated guesses when making decisions, solving problems, or answering questions. Because of its unique features, AI is often misunderstood and can carry risks, including security issues, mistrust, and potential harm. To manage these risks, AI needs strong engineering practices and proper regulation.

Most AI systems aren't built to last

AI systems are only as good as the software they run on. Unfortunately, 73% of AI and big data systems show structural quality issues–ranging from poor maintainability to missing controls–putting long-term performance at risk.

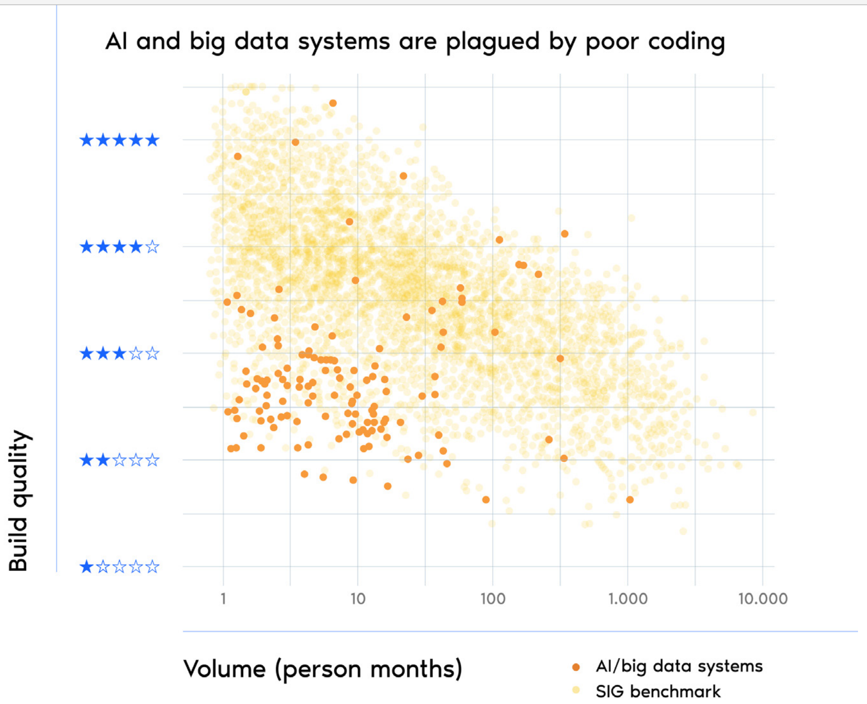

Our dataset of AI/big data systems was compiled by selecting systems that revolve around statistical analysis or machine learning, based on the technologies used (e.g. R and Tensorflow) and documentation.

Engineering robust, future-proof AI systems is still a relatively new field. Many organizations struggle to transition AI from experimental projects in the lab to scalable, secure, compliant, and maintainable real-world applications. Engineering challenges often stem from how AI engineers—such as data scientists—are traditionally managed and trained. Their focus is typically on generating insights and building models quickly, not on designing systems that are secure, reliable, testable, or easy to maintain.

SIG's data confirms this: AI and big data systems contain just **1.5% test code on average, compared to 43% in traditional systems**, making them harder to update and more prone to silent failure. Many systems are built with long, unfocused code blocks that combine multiple responsibilities, making them brittle, expensive to modify, and difficult to reuse.

Most AI systems are built in Python, which, as discussed in chapter 3, is a go-to language for data science and machine learning. While it is certainly possible to create high-quality software in Python, the way AI projects are typically developed means the resulting systems often carry high anticipated maintenance costs. Python also has a higher energy footprint compared to other languages, which can affect long-term sustainability. For a detailed comparison of Python's efficiency and maintainability versus other languages, see Chapter 3.

In the energy sector, where uptime and safety are critical, these risks compound quickly. To prevent AI becoming tomorrow's legacy, energy companies must treat it like critical infrastructure—starting with software engineering discipline from day one.

The path forward: Build AI like critical infrastructure

To realize long-term value from AI, energy companies must treat it as serious software rather than as an experiment. That means applying the same disciplined engineering practices used in core systems, from architecture and documentation to testing and governance.

SIG's data and client work suggest four essential steps:

- **Apply software engineering best practices** → Modular, maintainable, and well-documented code is essential for building AI that can evolve over time.
- **Bridge AI and software engineering teams** → Many AI systems are still developed in silos. Integrating software engineers into AI projects improves long-term stability and avoids creating legacy from day one.
- **Strengthen AI governance** → Organizations should define clear accountability and align AI development with policies for risk, compliance, and transparency.
- **Implement continuous validation** → AI systems degrade over time. Regular testing and retraining are critical to ensure safe, scalable performance.

Treating AI as a core capability—built to scale and governed with intent—is the only way to unlock its full potential in the energy sector.

Chapter 5: Legacy systems can slow energy down

Key findings:

- Architecture quality in the energy sector mirrors the industry average at around 4 stars. Most systems are change-ready, but not entirely future-proof.
- 4-star systems allow changes to be made **30% faster** than market average, while 2-star systems experience a **40% slowdown.**
- Legacy systems carry security and cost risks: **a 300% increase in expected risk costs**, and **twice the likelihood of a breach**

Modernizing the energy sector starts with software

The energy industry is undergoing rapid digital transformation. Companies are investing in smart grids, automation, AI, and data-driven operations to improve resilience and efficiency. According to the International Energy Agency (IEA), digital technologies will play a critical role in enabling decarbonization, resilience, and real-time grid optimization.

But this transformation is being held back by legacy infrastructure. According to IBM, 21% of utilities report making no progress toward grid modernization, and more than 40% of transmission and distribution investments are still being spent just to catch up. Recent academic studies also highlight that outdated infrastructures delay the integration of renewables, slow down analytics adoption, and raise operational risk–especially in organizations with lower digital maturity.

The challenge is clear: while digital ambitions are rising fast, many of the systems they rely on weren't built to support this pace of change. Without modern, adaptable architecture, energy companies risk falling behind in the transition to low-carbon, software-driven operations.

The architecture gap: where energy teams stand today

SIG's analysis shows architecture quality in energy is **on par with other sectors,** averaging around 4 stars. This gives teams a solid starting point, but modernization is still critical to keep pace with transformation.

The most modern architectures enable teams to make changes 30% faster than the industry baseline. Meanwhile, 2-star architecture systems are 40% slower to evolve.

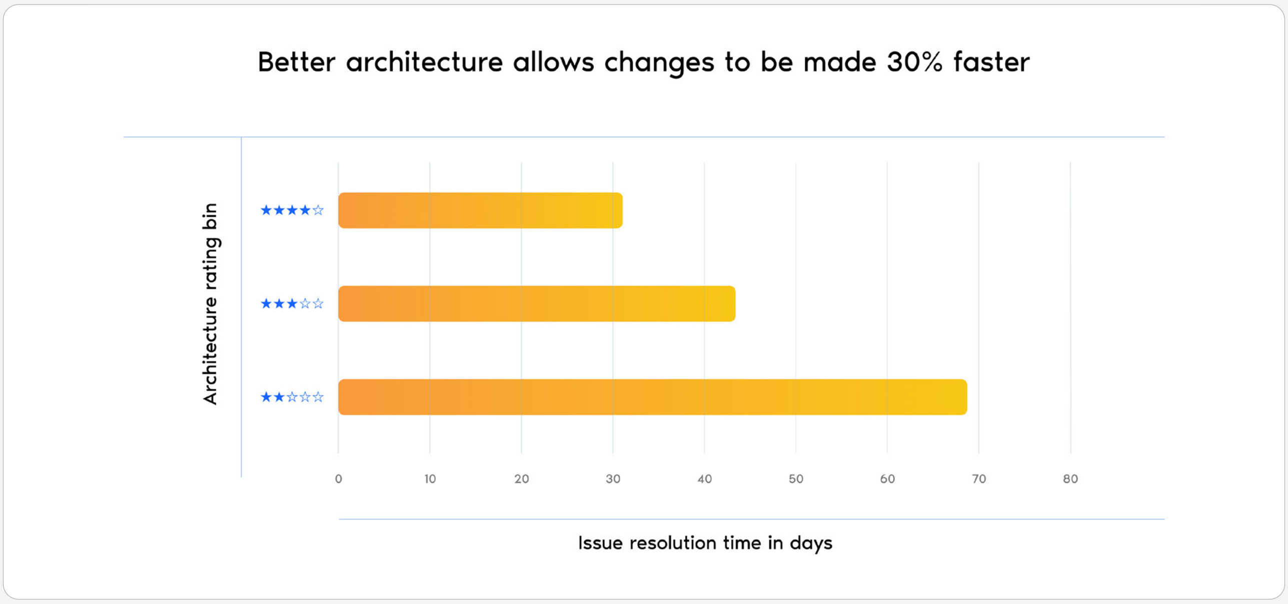
Better architecture allows changes to be made 30% faster

This diagram shows a correlation between architecture quality and issue resolution time. When it comes to architecture quality, it is, on average, 30% faster to make changes in a 4-star system than in a market-average system. Reversely, making changes to a 2-star system will be about 40% slower.

What makes good architecture?

Modern architecture is about design that enables flexibility, maintainability, and independence across teams. SIG's Architecture Quality framework identifies key characteristics of high-quality software architecture:

- **Clear component boundaries** → easy to isolate, understand, and maintain
- **Balanced component sizes** → avoids monoliths and fragmentation
- **Low coupling between components** → limits ripple effects when changes occur
- **High cohesion within components** → related logic stays grouped
- **Minimal code duplication** → less redundancy, fewer bugs
- **Efficient data access** → clear, fast paths to needed information
- **Standardized technology usage** → common, up-to-date tools and patterns
- **Distributed knowledge** → no single points of team failure
- **Clear communication paths** → interfaces are defined and predictable
- **Evolutionary flexibility** → architecture adapts as needs evolve
- **Scalable structure** → designed to handle future growth

These principles directly support cost-effective innovation and faster delivery. Both of which are critical in a sector under pressure to reduce costs while modernizing infrastructure.

The cost of doing nothing

The longer legacy systems remain in place, the harder and more expensive they become to replace. Teams become reluctant to touch fragile code, compounding technical debt over time.

Legacy systems also introduce risk. Organizations that rely heavily on legacy systems see a 300% increase in expected risk-related costs and are twice as likely to experience a major security incident.

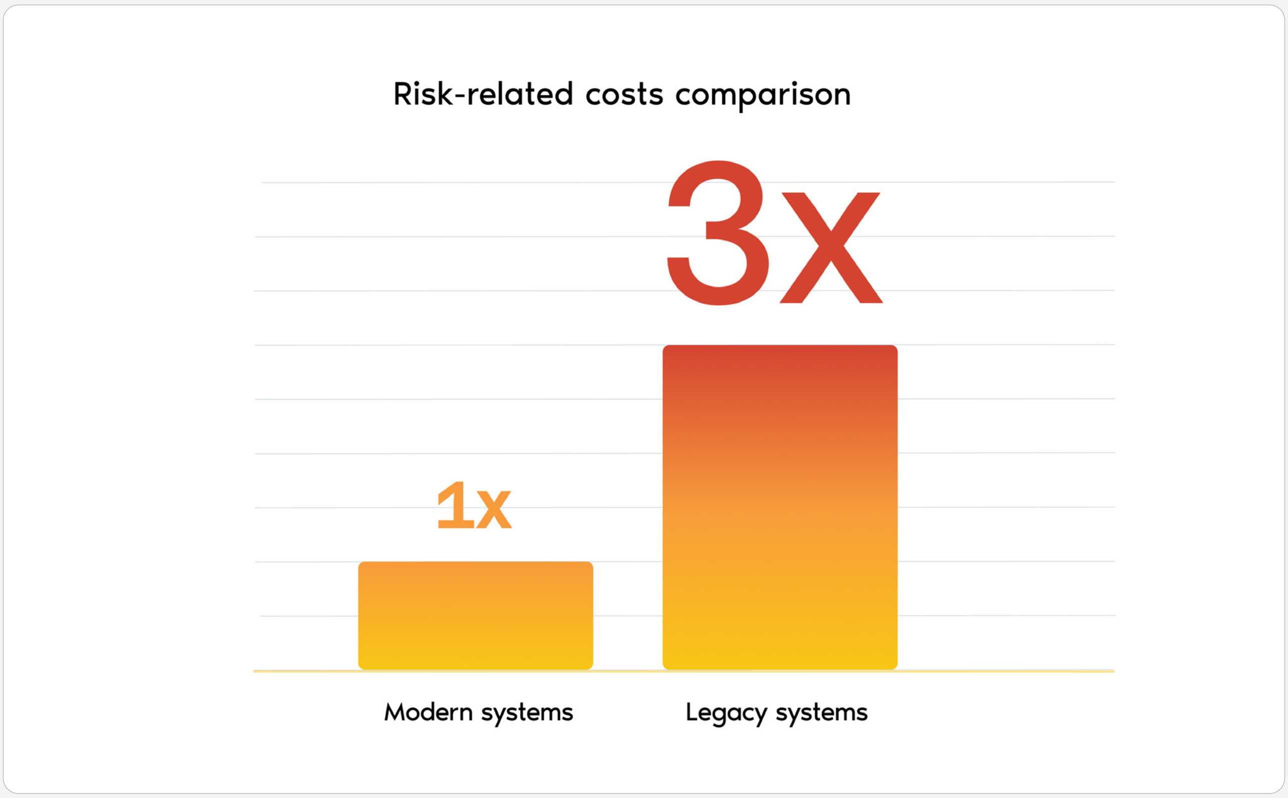

In the energy industry, software increasingly controls physical infrastructure. This means these risks are not theoretical. They're operational, regulatory, and reputational.

Chapter 6: Cloud readiness gives energy a competitive edge

Key findings:

- Energy systems outperform other industries on cloud readiness, averaging **3.9 stars** compared to a **cross-industry average of 3.1.**
- SIG's cloud readiness assessment measures six key dimensions: business impact, external interfacing, automation, technology usage, evolution, and knowledge distribution.
- Cloud-ready systems scale faster, cost less to operate, and are easier to evolve.

Why cloud readiness matters for energy

Cloud migration has become a strategic priority across industries, but in the energy sector, it's particularly urgent. From managing grid data and customer platforms to deploying AI-driven optimization, energy providers rely on digital infrastructure that must scale, evolve, and integrate quickly.

The potential impact is significant. McKinsey reports that cloud-powered AI, ML, and IoT accelerate nearly half of decarbonization initiatives—cutting costs by 2–10% and potentially abating 1.5 Gt CO_2e per year—even as total cloud adoption could unlock up to $3 trillion in global EBITDA by 2030.

To make this transition successfully and realize these benefits, energy systems must be "cloud-ready" i.e. technically and organizationally prepared to migrate to and operate efficiently in a cloud environment. This includes factors like platform-independent code, automation support, modular architecture, and clear knowledge distribution across teams.

Moving to the cloud brings several tangible benefits for energy firms. It helps optimize operations, reduce costs, and enhance agility by streamlining unstructured data management and enabling faster integration with partners and platforms. This operational flexibility is especially valuable in the energy sector, where demand can fluctuate rapidly.

Energy is ahead of the curve—but there's room to improve

SIG analysis shows that energy companies are in a strong starting position, with **an average cloud readiness score of 3.9 stars**, well above the **broader average of 3.1.** This reflects the sector's relatively modular systems and recent modernization efforts.

Cloud readiness score

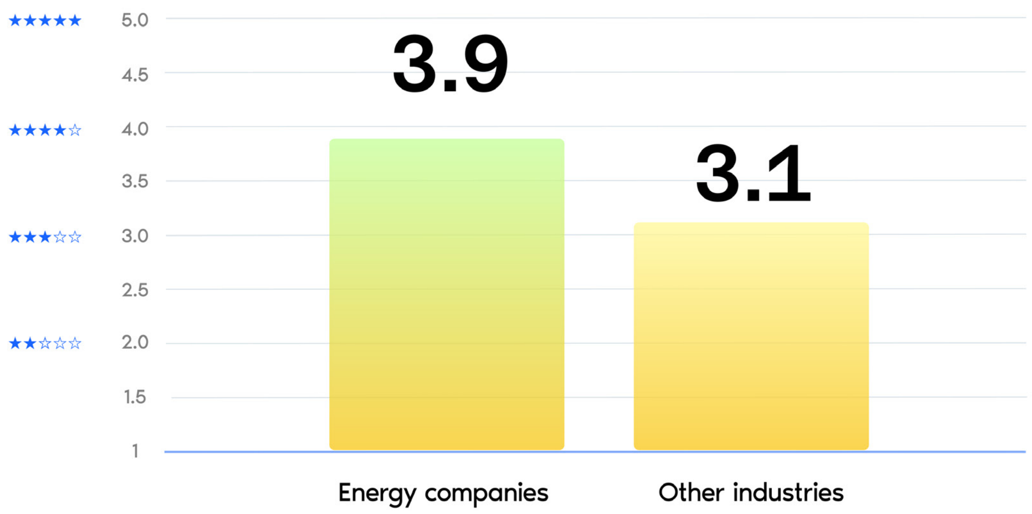

What we mean by cloud readiness

The SIG Cloud Readiness model assesses six critical areas:

- **Business impact** – *How much value would cloud migration deliver for this system?*
- **External interfacing** – *How cleanly does the system connect to other services?*
- **Automation** – *Are deployment and operations automated or manual?*
- **Technology usage** – *Are the tools and frameworks cloud-compatible?*
- **Evolution** – *How easily can the system be updated?*
- **Knowledge** – *Is the system well-documented and understood by multiple team members?*

But scores alone aren't enough. Cloud readiness is a spectrum. And in systems where readiness falls short, migration becomes costly, brittle, or delayed.

Cloud-ready systems move faster and cost less

Beyond technical readiness, cloud-capable systems offer real business value:

- **Faster deployment and evolution** – agility in delivering new features
- **Cost control** – scale resources up or down to match real demand
- **Improved reliability** – leverage cloud-native resilience and failover
- **Regulatory readiness** – support evolving compliance requirements with scalable infrastructure
- **Reduced vendor lock-in** – modular design supports portability across providers

Cloud readiness also supports broader goals like ESG targets, modernization strategies, and reducing total cost of ownership.

Don't migrate before you're ready

Modernizing systems for the cloud doesn't happen by accident. To stay competitive, energy companies must identify which systems are cloud-ready today, prioritize upgrades for those that aren't, and minimize risk during complex migrations. That means building future-facing software now—so it won't require rework two years from now.

In the energy sector uptime, cost control, and innovation all matter. Making cloud readiness a strategic enabler, rather than just a technical metric.

Chapter 7: Shared knowledge powering energy transformation

Key findings:

- 59% of energy companies meet SIG's recommended level of knowledge distribution. Far better than the cross-industry norm, where **70% of organizations fall short.**
- Poor knowledge sharing leads to higher costs, longer outages, and slower time to market
- Distributed knowledge supports compliance, accelerates learning, and preserves continuity in a rapidly evolving energy landscape.

Knowledge is infrastructure

The energy transition depends not only on renewable technology and digitalization, but also on the human systems that support them. Energy companies need skilled, distributed teams equipped to build, operate, and secure increasingly complex digital infrastructure.

Yet the sector faces a growing digital skills challenge. The IEA warns that while power companies increasingly require digital expertise, many lack active recruitment efforts to close the gap. Job turnover in energy remains high: one survey found that 60% of non-retirement attrition happens among employees aged 23–37, exacerbating the shortage of technical expertise.

This growing shortage of digital talent is a systems risk, not just a workforce issue. This "people-side" infrastructure is now critical to operational resilience. Without enough developers, engineers, and digital specialists, and without effective knowledge-sharing, energy firms risk being unable to maintain or evolve systems as complexity increases.

The IEA's World Energy Employment 2024 report confirms that labor shortages in digital and clean energy skills remain a significant barrier to scaling investments, particularly in areas like software engineering, AI integration, and smart systems deployment. Bridging this gap requires targeted hiring, internal upskilling, and knowledge transfer.

When key knowledge is siloed—when just one or two team members understand how a system works—it becomes a hidden point of failure.

Energy is ahead but vigilance is still needed

SIG analysis shows that **59% of energy companies meet the recommended level of knowledge distribution**, compared to just **30% across other industries.** That's a strong position, but not a guarantee. As systems grow and teams change, knowledge risks can creep in unnoticed.

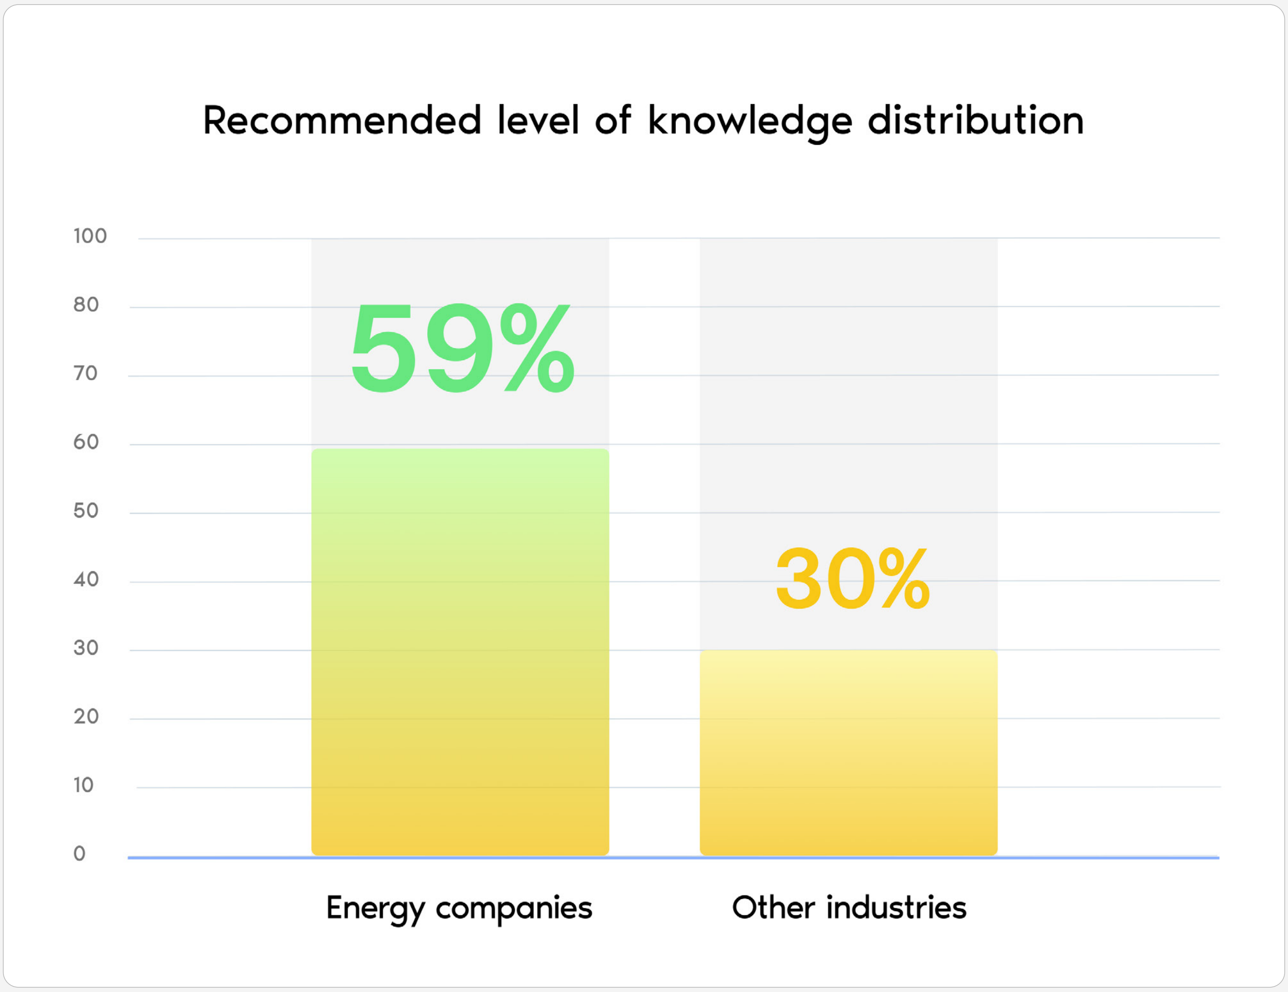

The risks of knowledge monopolies

Beyond technical readiness, cloud-capable systems offer real business value:

• **Higher operational costs**	→	New engineers struggle to navigate undocumented or complex legacy systems.
• **Longer recovery times**	→	When key experts leave, it takes longer to resolve issues and restore service.
• **Slower system evolution**	→	Poorly distributed knowledge limits agility, delaying critical innovations like smart grid deployment or AI integration.

For energy companies managing increasingly complex digital systems, these risks scale quickly—draining resources and threatening both reliability and innovation.

How knowledge monopolies form: The risk of uneven knowledge distribution

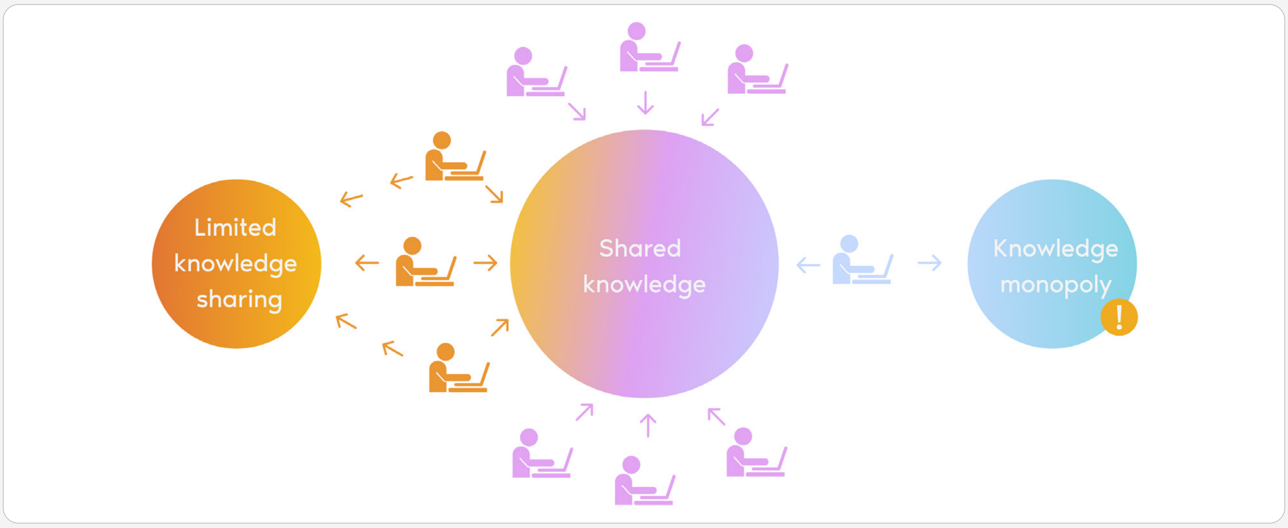

If you don't see it, you don't know it: The danger of neglected systems

When components aren't consistently updated or maintained, systems become harder to evolve, introducing security, cost, and compliance risks. In an industry where uptime is non-negotiable, these blind spots create unnecessary exposure.

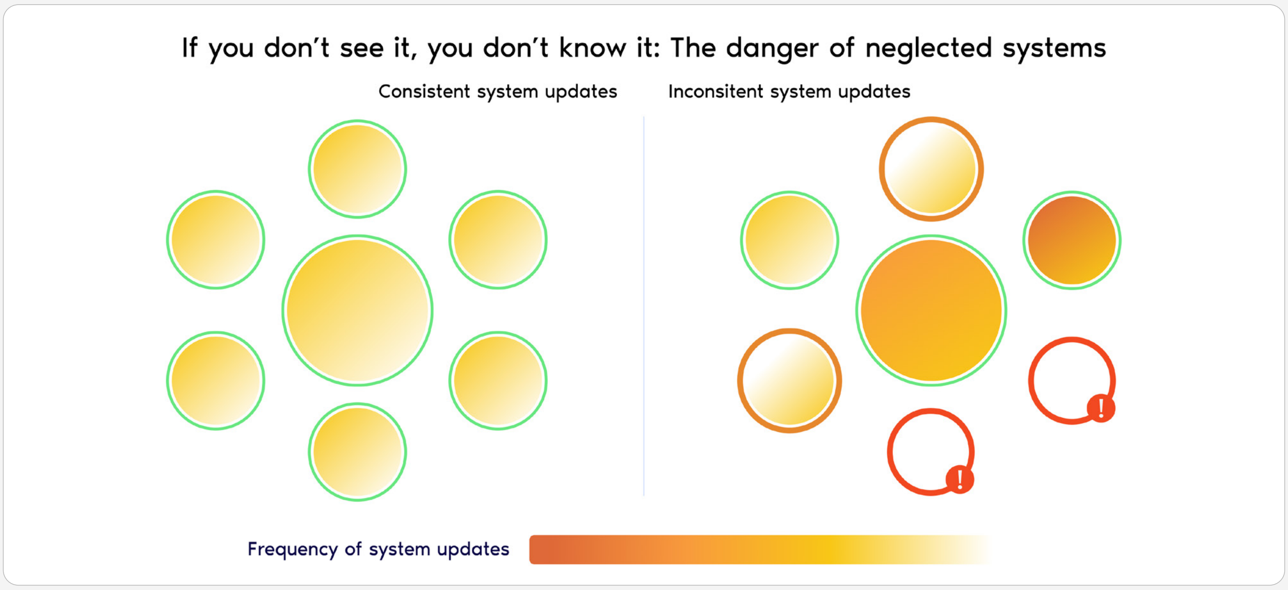
If you don't see it, you don't know it: The danger of neglected systems
Consistent system updates
Inconsitent system updates
!
!
Frequency of system updates

Prioritizing skills and knowledge distribution

To stay resilient, energy organizations must embed knowledge-sharing and upskilling into their digital transformation efforts. From predictive maintenance to cloud-native platforms, innovation depends on how well teams understand the systems they run.

Regular mentoring, documentation and training are essential tools for minimizing operational risk and enabling sustainable change.

Key takeaways & conclusion

- **Cybersecurity is critical to resilience.** → Cyberattacks on critical infrastructure are rising fast, so security must be embedded at the software level to protect operations.
- **Build quality drives performance and cost.** → Better software quality cuts downtime, lowers maintenance effort, and frees up resources for innovation.
- **Green IT is essential, not optional.** → Inefficient code wastes energy. Optimizing software can reduce emissions and operating costs while supporting climate targets.
- **AI offers value, but only with quality.** → 73% of AI systems have structural flaws. Without strong foundations, AI becomes harder to scale, maintain, and trust.
- **Modernization is a must.** → Legacy systems slow down change and introduce risk. Architecture quality directly impacts delivery speed and scalability.
- **Cloud readiness is a competitive edge.** → Cloud-ready systems are easier to scale, adapt, and secure. Energy leads the pack–but must continue to close the gaps.
- **Knowledge is infrastructure.** → Poor knowledge sharing raises costs and risk. Distributed expertise helps energy companies innovate, comply, and recover faster.

The energy leaders of tomorrow prioritize software quality

Energy providers that take a structured, data-driven approach to software quality will lead the transition to a more secure, efficient, and sustainable future.

At Software Improvement Group (SIG), we help energy organizations assess, benchmark, and optimize their software portfolios so they're ready for what's next.

Want to see how your systems compare?
Contact us today to take control of your software landscape.

Written by Software Improvement Group

Software Improvement Group (SIG) leads in traditional and AI software quality assurance. Empowering organizations to become more resilient and agile by guiding them to enhance their software quality and security through deep source code analysis and tailored, strategic advice.

Sigrid® - its software assurance platform - leverages the world's largest database containing over 300 billion lines of code across more than 20,000 systems and 300+ technologies and intelligently recommends the most crucial initiatives for organizations. SIG complies with multiple ISO/IEC standards, including ISO/IEC 27001 and 17025, and has co-developed ISO/IEC 5338, the new global standard for AI lifecycle management.

SIG was founded in 2000 and has offices in New York, Copenhagen, Brussels, and Frankfurt, and is headquartered in Amsterdam.

Sigrid®, together with expert software engineering consultants, and over 25 years of industry-leading research, position SIG as the foremost authority on software excellence.

Trusted by

alliander	EDSN energiedata voor iedereen	essent	Eneco
TenneT	STEDIN NET	ENEXIS ENERGIE IN GOEDE BANEN	gasunie

Ensure your software is always two steps ahead

Modernize your grid, secure your IT, and stay in control.

Discover SIG for Power & Utilities

Energy signals 2025

Software Improvement Group