

RESEARCH REPORT

Energy reliability in the economics of industrial location

Industrial energy costs depend on reliable delivery, connection timing, and the production process. Four official records and original cost scenarios show how interruptions, backup, and delayed operation can change a location comparison.

Information cutoff: 2026-10-27

<https://aesenate.org/research/energy-reliability-in-the-economics-of-industrial-location/>

The relevant energy cost for an industrial location is the cost of running the intended production process to the required standard. A quoted electricity tariff captures only part of that cost. Network access, the timing of the connection, interruptions, voltage quality, backup resources, and exposure to variable prices can change the economics materially. Their importance depends on the process: an interruption that is manageable for an office can destroy a batch, delay a delivery, or damage equipment in another activity.

This report examines official evidence from the United States, Germany, India, and South Africa and develops an original delivered-energy cost comparison. The evidence describes different parts of the electricity system. It does not produce a harmonized ranking of national industrial reliability. The numerical scenarios are expressly hypothetical and compare two unnamed sites; no tariff, outage assumption, or cost estimate in those scenarios is attributed to one of the four economies.

The central conclusion is that reliability should enter location appraisal through the production losses and mitigation costs a firm actually faces. Connection timing should be evaluated separately because it affects when the entire investment can earn a return. Price risk then needs to be matched with the firm's load and contract. A low average tariff, a low system energy deficit, and a favorable interruption index each provide useful evidence, but none establishes all three conditions.

What four official records show

The U.S. Energy Information Administration reports that electricity customers experienced an average of 11 hours of interruptions in 2024. Major events accounted for 80 percent of those hours. Its December 1, 2025 analysis explains both interruption duration and frequency and shows substantial differences across states. The national average is therefore a description of a distribution-system record, not an estimate of the outage exposure of a particular industrial feeder or facility. [EIA, 2024 electricity interruptions](#).

Germany's Bundesnetzagentur reports 11.7 minutes of average unavailability for 2024 in a release dated October 9, 2025. The published SAIDI measure covers relevant unplanned interruptions and excludes events classified as force majeure. Its reporting includes interruptions longer than three minutes. Comparing that number directly with the American all-event figure would ignore a material difference in coverage. The figures should be retained with their definitions rather than converted into an apparently precise relative reliability score. [Bundesnetzagentur, 2024 interruption statistics](#).

India's Ministry of Power, in a parliamentary reply published February 5, 2026, reports an all-India energy gap of 0.1 percent in financial year 2024-25. Its table compares energy required and supplied and separately reports the peak position. These are system supply measures. They are not a customer-level record of the number, length, or voltage characteristics of interruptions at an industrial site. [Ministry of Power, power demand and generation capacity, Annex I](#).

Eskom's 2025 annual financial statements report 175 hours during which load shedding and load curtailment were implemented, against 6,367 in the prior reporting year. The measure records system operations rather than the outage duration experienced by every customer. The same operational discussion reports improved generation availability. That improvement is relevant evidence, but it does not imply that every transmission line, distribution connection, or industrial process had an identical experience. [Eskom, 2025 financial statements, operational performance](#).

Each source identifies a different measurement boundary: customer interruptions, selected unplanned interruptions, energy balance, and system load-shedding operations. The differences are economically important. A generation system can have adequate annual energy while a local network connection remains constrained. A customer can have few interruptions but face a long wait for additional capacity. National evidence should guide further inquiry, not substitute for the terms and engineering assessment of the proposed site.

The unit of service is usable production power

A manufacturing process requires power at a specified location, voltage, quality, and time. An annual quantity of megawatt-hours does not establish that the system can deliver the required megawatts during a critical production interval. A site with a flat load differs from one with short peaks. A process that can shift its operating hours differs from one that must run continuously. The service specification should be established before electricity offers are compared.

The distinction between energy and capacity is particularly important for backup design. A battery may hold enough energy for a short interruption but lack the power rating to start a large motor. A generator may supply the required power but depend on fuel delivery during the same event that disrupts roads. A second electrical feed can improve resilience only to the extent that it avoids the failure affecting the first. These are engineering conditions with direct economic consequences.

Power quality can matter even when an interruption statistic records little or no lost time. A short voltage disturbance may require a production reset or damage a sensitive product. The economic assessment should

identify the tolerance of the actual equipment, the quality protection required, and the evidence available from the network operator. A broad duration index cannot establish the frequency or cost of every disturbance relevant to that process.

The timing of an event matters as well. An interruption during planned maintenance can have a lower incremental cost than an identical interruption during a continuous production run. A scheduled curtailment can permit advance adjustment but may coincide with high-value output or delivery commitments. An unscheduled event can create additional restart, safety, or spoilage costs. Expected outage hours alone do not capture this distribution of consequences.

A reliable cost comparison therefore starts with a load profile and a production profile. It specifies required power, annual consumption, the cost of curtailed production, the scope for rescheduling, and any minimum safe operating supply. The appraisal then examines the system and contract capable of meeting those requirements. Starting with the cheapest published tariff and adding an arbitrary reliability premium reverses the order of the economic question.

An original comparison of two sites

Consider an illustrative facility requiring 10,000 megawatt-hours each year at its planned output. At site A, the grid price is 80 monetary units per megawatt-hour and interruptions would prevent delivery of 50 megawatt-hours. At site B, the grid price is 90 and no interruption is assumed in the base scenario. These prices and service assumptions are hypothetical, not quotations or national averages. Fixed network charges are assumed equal and omitted so the reliability mechanism remains visible.

Without backup, site A buys 9,950 megawatt-hours at a cost of 796,000. The 50 megawatt-hours not delivered are associated with lost contribution of 800 per megawatt-hour, or 40,000. Four interruptions also produce separate restart and spoilage costs of 5,000 each, or 20,000. The contribution-loss assumption excludes those additional costs to avoid double counting. Total annual energy-related economic cost is therefore 856,000.

Now give site A backup that supplies all 50 missing megawatt-hours. Assume its power rating, transfer arrangements, and fuel availability are sufficient to prevent the modeled production and restart losses. Backup variable cost is 250 per megawatt-hour and annual fixed ownership and maintenance cost is 12,000. Grid cost remains 796,000; backup variable cost is 12,500; total cost is 820,500. The backup's full assumed capability is a condition of the result, not an engineering finding.

Hypothetical annual cost	Site A, no backup	Site A, sufficient backup	Site B, base scenario
Grid electricity	796,000	796,000	900,000
Backup variable cost	0	12,500	0
Annual backup fixed cost	0	12,000	0

Hypothetical annual cost	Site A, no backup	Site A, sufficient backup	Site B, base scenario
Lost production contribution	40,000	0	0
Restart and spoilage	20,000	0	0
Total	856,000	820,500	900,000

The comparison demonstrates two different decisions. Backup lowers site A’s modeled annual cost by 35,500. Site A with backup is 79,500 cheaper than site B under the stated assumptions. Neither result follows from the tariff difference alone. A comparison based only on purchased grid expenditure would omit lost production, while one adding backup expenditure and the same avoided production loss would incorrectly count both alternatives together.

Dividing by the common requirement of 10,000 megawatt-hours gives 85.60, 82.05, and 90 per required megawatt-hour. These are economic-cost equivalents for a common planned production requirement. They are not billed electricity prices, and the first case does not actually deliver the full quantity of power. That label is necessary to keep a useful comparison from implying that undelivered electricity was physically consumed.

Connection delay can reverse the first-year result

Annual operating cost does not settle the location decision if one site can start later. Suppose site A with backup faces an additional three-month connection delay relative to site B. Lost production contribution during that interval is assumed to be 150,000 per month, giving a one-time delay cost of 450,000. The figure represents a stated opportunity-cost assumption, not a forecast of revenue, and excludes costs already counted elsewhere.

Adding that delay cost to a full-year steady-state comparison raises site A’s illustrative first-year-equivalent burden to 1,270,500, compared with 900,000 at site B. This combination is a simplified appraisal device, not a calendar-year income statement: the operating costs describe a full operating year, while the delay is a separate initial penalty. A detailed project model would put expenditure and output on actual dates and include the operating costs avoided during the delay.

The annual saving of 79,500 would recover a 450,000 initial penalty after about 5.66 years on an undiscounted basis. With ten year-end annual savings and an 8 percent discount rate, the present value factor is the sum of 1 divided by 1.08 raised to each year from one through ten, or 6.71008. Savings have a present value of approximately 533,451; subtracting the initial penalty leaves approximately 83,451.

That narrow positive result is sensitive to the horizon, discount rate, and persistence of the operating advantage. It also omits differences in taxes, labor, logistics, emissions compliance, land, and terminal values. The exercise therefore cannot select an actual location. It demonstrates that a connection delay can dominate

short-term savings and that the importance of a tariff advantage depends on how long the project operates after the delay.

The contractual status of a connection should consequently be examined directly. An application, a preliminary offer, an accepted agreement, construction progress, energization, and firm operating capacity are different stages. A location announcement based on an unconfirmed capacity date exposes the project to a different risk from a site with an energized connection. The expected date should be tied to the network works and approvals required, rather than to a promotional timetable.

Germany's network-expansion monitoring illustrates this stage distinction at a system level. It records planning and construction progress for specified transmission projects and explicitly excludes distribution-line projects. Those records can help establish wider network development, but they do not verify the date or capacity of a particular industrial distribution connection. [Bundesnetzagentur, network expansion monitoring, second quarter 2026](#).

Duration and frequency change the mitigation decision

A given total of annual outage hours can require different backup arrangements. Ten short events may cause more restarts than one long event, while one long event can exceed the energy duration of a battery. The facility's cost function should therefore include frequency, duration, timing, and restoration sequence. An estimate based on average hours multiplied by average lost output can understate the effect where costs rise sharply after a particular threshold.

The illustrative site's restart cost was entered separately for this reason. If the same 50 megawatt-hours of lost supply occurred in one event rather than four, the modeled restart cost would fall from 20,000 to 5,000, holding all other assumptions constant. The physical energy shortfall is unchanged. Conversely, if an event requires a long cleaning or recalibration process, the production loss can continue after grid power has returned. The relevant interruption to production can exceed the interruption to electrical supply.

Mitigation should be assessed against this event distribution. A small uninterruptible supply might protect control systems and permit safe shutdown without sustaining full production. A battery might bridge the transfer to a generator. Flexible scheduling might avoid a predictable curtailment at lower cost than additional equipment. A firm should compare the combinations that meet the process requirement rather than assume that complete self-generation is always necessary.

Expected cost is only one decision criterion. A rare severe event can threaten a delivery contract, customer relationship, or the firm's ability to continue operating. Those consequences can justify protection beyond the option with the lowest average annual cost. The analysis should identify the loss distribution and risk-bearing capacity rather than hide the judgment inside an unexplained premium. Insurance can transfer part of a financial loss but may not restore the customer or the physical production schedule.

Correlation is relevant to a second site or backup supplier. Two facilities exposed to the same weather system, transmission corridor, fuel route, or equipment supplier can fail together. Their combined capacity may be less

resilient than their separate availability records suggest. Geographic diversification has value when it changes the underlying failure exposure, not merely when it adds another address to the production network.

Testing the break-even assumptions

The illustrative backup decision can be expressed as a threshold rather than a single favorable case. With four restart events, the avoided restart cost is 20,000. Backup variable cost is 12,500 and its annual fixed cost is 12,000, giving a total backup cost of 24,500. Backup therefore requires at least 4,500 of additional avoided production loss to break even. Across 50 missing megawatt-hours, that corresponds to a contribution loss of 90 per megawatt-hour. The original assumption of 800 is well above that threshold.

If the same energy shortfall instead occurs in one event, avoided restart cost is only 5,000. The required avoided production loss rises to 19,500, or 390 per missing megawatt-hour. The energy quantity and backup price have not changed; the value of avoiding repeated starts has. This sensitivity shows why a reliability appraisal should use information on event frequency rather than infer it from annual unserved energy.

The maximum annual fixed backup cost can also be calculated. In the original four-event case, avoided production and restart losses total 60,000. Subtracting variable backup cost of 12,500 leaves 47,500 available for annual fixed cost before the backup option ceases to reduce the modeled economic cost. That threshold should be compared with a complete ownership estimate, including replacement, maintenance, testing, and the opportunity cost of committed capital. A purchase price without its operating life is not the same input.

Those thresholds still depend on effective protection. If transfer time causes the production process to stop, backup may avoid only part of the assumed restart loss. If fuel is unavailable during a long event, the expected energy supplied is lower. The appraisal should revise the avoided-loss calculation accordingly rather than retain full protection and merely add a general risk allowance. The mechanism determining the residual loss should remain identifiable.

Commercial evidence before commitment

An industrial buyer needs evidence from more than a price proposal. The connection offer should identify firm and interruptible capacity, required network works, dependencies, and the conditions attached to the proposed date. The supply contract should establish which charges are fixed, which vary, and what happens during curtailment. Equipment specifications should establish the production load and the backup performance required. These documents answer different questions and should be reconciled into one operating scenario.

A delay allocation in a contract can compensate the buyer without making the facility productive sooner. The value of that compensation depends on enforceability, exclusions, limits, and the counterparty's ability to pay. It should not be treated as a physical substitute for a connection. Similarly, an insurance payment can offset some financial losses while leaving delivery commitments unresolved. Commercial protection and physical resilience are complementary margins that should be priced separately.

For a staged investment, the firm can also compare a smaller initial load with later expansion. An earlier limited connection may allow training, validation, and initial production while broader network work proceeds. That option has value only if the partial operation serves a genuine commercial or technical purpose. Its costs include duplication, later interruption, and the risk that expansion never becomes available. The appropriate comparison is the present value of credible operating paths, rather than the largest capacity shown in an announcement.

Price volatility and the load profile

A simple average electricity price is not necessarily the price paid by an industrial user. A facility consuming more during high-price intervals faces a load-weighted price. Demand charges, network charges, balancing obligations, taxes, and contract terms can add further differences. A comparison of industrial tariffs should specify voltage level, load factor, consumption band, and the treatment of those charges before interpreting the resulting gap.

Contracts redistribute price risk. A fixed-price agreement can reduce short-term exposure while carrying a price for that protection or conditions on volume. A variable-price contract can benefit from low-price periods while exposing the firm to expensive intervals. A long-term renewable agreement may require additional arrangements to match intermittent output with the facility's continuous load. The contract's quoted generation price does not alone establish the delivered cost of reliable power.

The relevant hedge can differ from the risk being described. An instrument linked to an average market price may leave exposure to the difference between that price and the facility's local or hourly cost. Volume risk remains if production changes. A physical outage can coincide with a price spike and a need to purchase replacement power. Contract evaluation should therefore distinguish price, volume, location, and physical-delivery risk.

Flexibility can be an economic asset. A process that can shift demand without losing output may avoid high-price intervals or provide a service to the system. The cost of flexibility includes inventory, scheduling, labor arrangements, and any equipment wear. Its benefit should be evaluated against those costs and the reliability of the price signal, rather than treat every reduction in peak consumption as a free efficiency gain.

For a multinational project, currency exposure adds another layer. An apparently cheaper tariff in one currency can change relative to export revenue or foreign-currency debt. Indexation can also move the local tariff with fuel or exchange rates. A location comparison should specify the currency of its costs and revenues and avoid presenting a nominal tariff converted at one spot rate as a permanent operating advantage.

Reliability, efficiency, and environmental requirements

Energy efficiency can reduce both ordinary expenditure and the backup capacity needed for a given output. It can also require capital and maintenance. An efficient process that is more sensitive to disturbances may need different protection from an older process. Evaluating efficiency and reliability together helps determine the least-cost production system, while keeping their distinct benefits visible.

Environmental and permitting requirements belong in the same cost boundary where they affect the feasible backup option. A diesel generator, gas plant, battery, or contracted alternative can face different operating limits, local conditions, and replacement needs. This report does not assume that the backup in the numerical example is permissible at any actual site. A real appraisal must establish the applicable requirements before treating an assumed option as available.

The treatment of emissions should be consistent across locations. A tariff comparison that excludes a required carbon or emissions cost in one site and includes it in another misstates the commercial choice. A broader public-policy comparison may also examine costs not priced in the private contract. The private financial result and the social evaluation can differ; they should be reported separately rather than blended into an unexplained adjusted tariff.

A lower electricity cost can also be offset by higher consumption per unit of accepted output. Process yield, equipment efficiency, cooling needs, and product mix affect that requirement. Comparing price per megawatt-hour is therefore an intermediate step. The final operating comparison should connect reliable delivered electricity to output meeting the same specification, including scrap, rework, and any load used while production is interrupted.

Interpreting the four cases without false equivalence

The American record directs attention to the contribution of major events and the variation hidden by a national average. The German record demonstrates that a selected interruption index needs its event exclusions attached. India's energy-balance record provides evidence of the system supply position while leaving site-level distribution performance to be established. Eskom's operations record documents changes in system load shedding without specifying the complete experience of each customer.

These are separate empirical questions rather than conflicting descriptions of one universal reliability variable. A harmonized comparison would need a common customer class, event threshold, exclusion rule, period, and location detail. Even then, an average customer index would need to be translated into the load and failure sensitivity of the intended facility. The most useful international comparison is therefore explicit about which stage is comparable and which requires additional evidence.

Institutional causes should remain separate from measurement. A transmission bottleneck, distribution fault, generation shortage, and delayed connection can create similar production losses through different mechanisms. The remedy may lie in maintenance, network expansion, contracting, backup, or project scheduling. Assigning all four to a single category of national energy risk would make the location analysis less precise and the policy response less well targeted.

Improvement over time should also be interpreted carefully. A better annual record can reflect investment, maintenance, weather, demand, or other conditions. The record is useful evidence without establishing that the same improvement will continue or that one intervention caused it. A forward-looking project should test whether the network works, fuel arrangements, and operating resources supporting that improvement remain credible over its investment horizon.

Conclusion

Industrial location appraisal should compare usable power at the required time and quality, including mitigation and lost production, rather than rely on a tariff alone. The original scenarios show that backup can be economical despite raising purchased-energy expenditure, while connection delay can absorb years of an operating advantage. The conclusion depends on disclosed assumptions about the process, event distribution, and project horizon.

Official evidence from the four economies is valuable precisely when its measurement boundary is retained. Interruption duration, energy balance, and load-shedding hours should not be combined into a false national league table. Connecting those observations to a specific site's engineering and commercial requirements produces a more defensible account of reliability and a clearer investment decision.

For the German network and investment relationship, see [Infrastructure bottlenecks and German capital formation](#). For the distinction between expenditure and usable productive services, see [The composition of American economic growth](#).

Source addresses

The links below retain the source addresses cited in the analysis. References should be read within their stated periods and scope.

<https://www.eia.gov/todayinenergy/detail.php?id=66744>

https://www.bundesnetzagentur.de/SharedDocs/Pressemitteilungen/DE/2025/20251009_Saidi_Strom.html?nn=660040

<https://www.pib.gov.in/PressReleasePage.aspx?PRID=2223710&lang=1®=1>

<https://www.eskom.co.za/wp-content/uploads/2025/09/Eskom-annual-financial-statements-2025.pdf>

<https://www.netzausbau.de/Daten/monitoring/de.html>