WHITE PAPER 2

From N-1 to Virtual N-1

Extending the Risk-Based Framework for Grid Capacity

Svendsen* · Nyiredy · Milenkovic****

May 2026

Contact: robert.nyiredy@siemens.com | arne-brufladt.svendsen@siemens.com

> ***“The N-1 criterion has served us well for sixty years. What Virtual N-1 does is not replace it — it makes it legible, differentiated, and manageable in a world that N-1 was never designed for.”***

This paper builds on: Svendsen, Nyiredy & Milenkovic (2026). Unlocking Grid Capacity Through Risk-Based Operation. White Paper 1.

Table of contents

Glossary of Acronyms

Acronym	Definition
ACER	Agency for the Cooperation of Energy Regulators
CACM	Capacity Allocation and Congestion Management (EU regulation)
CENS	Cost of Energy Not Supplied
DSA	Dynamic Security Assessment
EENS	Expected Energy Not Supplied
ENTSO-E	European Network of Transmission System Operators for Electricity
GET	Grid Enhancing Technologies
IEA	International Energy Agency
N-1	Single-contingency operational criterion
N-X	Multi-contingency state (X elements out of service)
PRA	Probabilistic Risk Assessment
SOGL	System Operation Guideline (EU network code)
TSO	Transmission System Operator
UPS	Uninterruptible Power Supply

Source Attribution

Throughout this paper, a distinction is made between:

- “IEA estimates” referring to quantitative values and findings directly reported in IEA (2026).
- “Authors’ interpretation” referring to system-level interpretations derived from these findings.
- “Operational experience” referring to results observed in practical implementations.

Unless otherwise stated, percentage-based estimates (e.g. 15–30%, 10–20%) represent authors’ interpretations rather than directly reported IEA results. Capacity uplift estimates are indicative and require validation through system-specific analysis before application to planning decisions.

This paper forms the second in a series. White Paper 1 (Svendsen, Nyiredy & Milenkovic, 2026) presents the three-layer operational framework for PRA, DSA, and Grid Enhancing Technologies. This paper extends that framework by introducing Virtual N-1 as the third and most structurally transformative stage of implementation.

Executive Summary

White Paper 1 established that 1,200–1,600 GW of additional global grid capacity could be unlocked through operational measures applied to existing infrastructure (IEA, 2026). It proposed a three-layer framework — a Decision Layer (PRA/DSA), a Physical Layer (Grid Enhancing Technologies), and a Security Layer (flexibility and markets) — to deliver an indicative 15–30% increase in utilisable grid capacity. That argument is complete on its own terms.

This paper takes the next step. It addresses a question that the three-layer framework takes as given but does not resolve: what is the N-1 security criterion, and is it the right foundation for the grid of the future? The argument proceeds in four steps:

- N-1 is not a safety target. It is a practical heuristic developed in the 1950s for a system with limited computational capability, homogeneous load, and predictable generation. It defines a boundary without describing what lies on either side of it — no probability, no consequence, no cost.
- Quantified N-1 (Stages 1–2 of the three-layer framework) makes the implicit N-1 margin explicit. By applying PRA and DSA in real time, operators can distinguish states where the margin is genuinely needed from states where it is excessive — unlocking 15–30% of utilisable capacity without changing the criterion itself.
- Virtual N-1 (Stage 3) goes further. It introduces conditional connection as a grid operating principle: large, flexible loads connect under contractual terms that accept interruptibility in exchange for earlier, faster, or cheaper grid access. The load itself becomes the active safety mechanism, maintaining effective N-1 security for all priority load through active consequence management.
- The incremental capacity potential from Virtual N-1 is estimated at an additional 10–20% of utilisable grid capacity. Combined, the full three-stage implementation indicates a 20–45% increase in utilisable capacity relative to today's deterministic N-1 baseline (authors' interpretation).

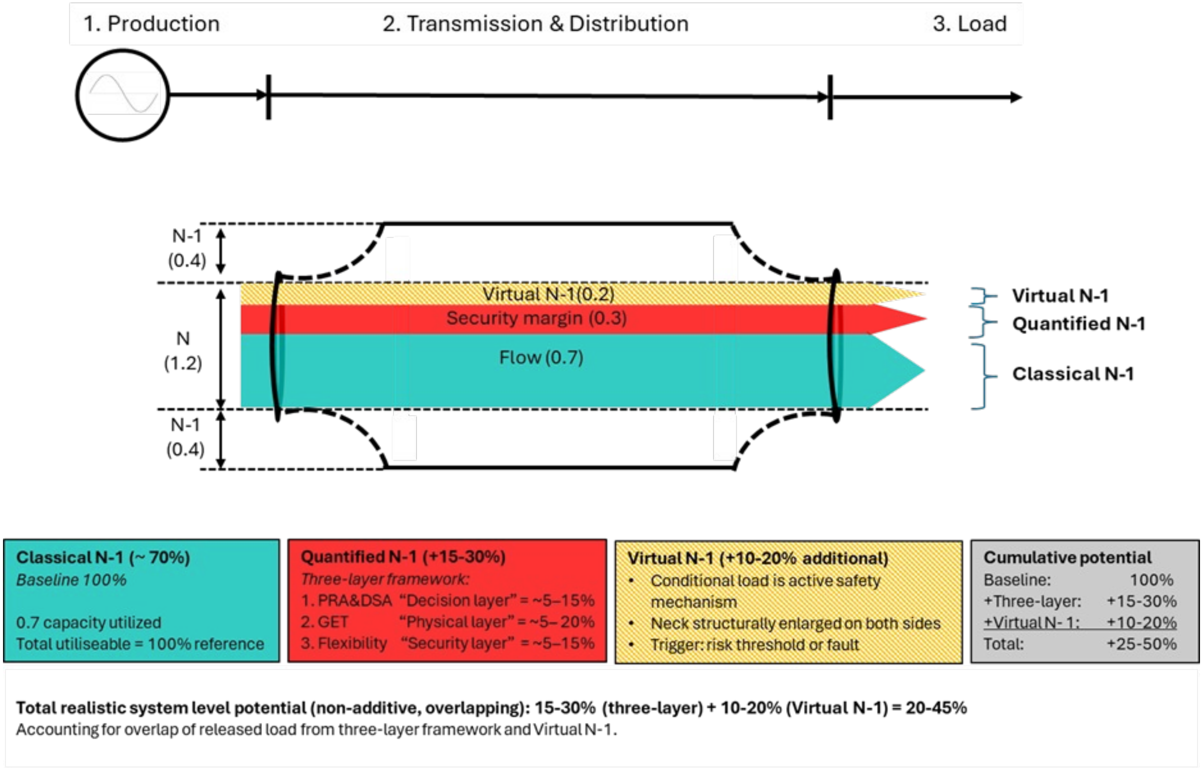

Classical N-1 (~ 70%)
Baseline 100%

0.7 capacity utilized
Total utiliseable = 100% reference

Quantified N-1 (+15-30%)
Three-layer framework:
1. PRA&DSA "Decision layer" = ~5–15%
2. GET "Physical layer" = ~5– 20%
3. Flexibility "Security layer" = ~5–15%

Virtual N-1 (+10-20% additional)
- Conditional load is active safety mechanism
- Neck structurally enlarged on both sides
- Trigger: risk threshold or fault

Cumulative potential

Baseline:	100%
+Three-layer:	+15-30%
+Virtual N- 1:	+10-20%
Total:	+25-50%

Total realistic system level potential (non-additive, overlapping): 15-30% (three-layer) + 10-20% (Virtual N-1) = 20-45%
Accounting for overlap of released load from three-layer framework and Virtual N-1.

Figure 1

> ***Virtual N-1 does not reduce system security. It redistributes the consequence of potential failure — from an uncontrolled, uniform outcome affecting all consumers, to a controlled, contractually defined outcome affecting only loads that have accepted, and been compensated for, conditional connection terms.***

The technologies and methodologies required are available. The operational experience exists — TenneT's non-firm contracts in the Netherlands alone unlocked 9.1 GW equivalent to 40% of national peak demand (Ember, 2026). The remaining work is definitional, regulatory, and commercial. It needs to start now, because the loads seeking Virtual N-1 connection are already in the queue.

1. Context: Building on the Three-Layer Framework

The electricity grid is now the primary constraint on the energy transition. Generation capacity continues to expand; new demand from electrification, data centres, and AI infrastructure is growing faster than anticipated; and a consistent finding across international assessments — including the IEA Electricity 2026 report — is that grids, not generation, are the binding bottleneck (IEA, 2026).

White Paper 1 proposed a structured response: a three-layer framework for risk-based grid operation, combining Probabilistic Risk Assessment (PRA) and Dynamic Security Assessment (DSA) as the Decision Layer, Grid Enhancing Technologies as the Physical Layer, and flexibility mechanisms as the Security Layer (Svendsen, Nyiredy & Milenkovic, 2026). Together, these deliver an indicative 15–30% increase in utilisable grid capacity by making the implicit N-1 security margin explicit and manageable.

This paper addresses the next level of the argument. The three-layer framework optimises operation within the existing N-1 criterion. It does not question the criterion itself. But to understand why 15–30% is achievable — and what the path to more looks like — we need to understand what N-1 actually is, and why it is structurally inadequate in today's system.

White Paper 1	White Paper 2 (This Document)
Baseline: N-1 as a given	Examines what N-1 is and is not
Proposes three-layer operational framework	Extends framework to Stage 3: Virtual N-1
Unlocks implicit margin (15–30%)	Expands effective capacity boundary (+10–20%)
Technical & regulatory focus	Conceptual, analytical & regulatory focus
Intended for technical and policy audiences	Intended for policy, regulatory & commercial audiences

Table 1

2. What a Power System Actually Does

The conventional description of a power system — a technical infrastructure for the production, transmission, and distribution of electrical energy — is accurate but misleading. It describes the physical structure while concealing the operational reality.

> *A more precise formulation: A power system is a real-time coordination system for the simultaneous balancing of energy, cost, and risk under continuous uncertainty.*

The system does not simply deliver power. It manages the consequences of failure continuously. Production varies. Consumption varies. Components fail. In this environment, the power system's actual task is not to prevent all failures — that is physically impossible — but to ensure that the consequences of failures remain acceptable to society.

This reframing matters for everything that follows. If the system's purpose is consequence management under uncertainty, then the question of how we define and enforce operating criteria is a question about how we quantify and allocate risk — not just how we design infrastructure. This is the conceptual foundation on which the three-layer framework rests (Svendsen, Nyiredy & Milenkovic, 2026), and on which Virtual N-1 builds.

3. N-1: A Historical Heuristic, Not a Safety Target

3.1 What N-1 Says

The N-1 criterion is the foundational design and operating principle of transmission systems worldwide. It states that the system must be able to withstand the loss of any single component — a line, transformer, or generator — without loss of supply to any load. It is applied deterministically: if any credible single contingency would cause a violation, the operating state is deemed insecure and must be changed. This rule has provided the basis for transmission planning, real-time operation, and regulatory compliance for over sixty years, and is embedded in grid codes and planning standards worldwide, including the ENTSO-E System Operation Guideline (ENTSO-E, 2023).

3.2 What N-1 Does Not Say

What the N-1 criterion does not do is at least as important as what it does:

- It says nothing about the probability that any given contingency will occur.
- It says nothing about the consequence of the contingency if it does occur.
- It says nothing about the cost — to the system, to consumers, or to society — of maintaining compliance.
- It treats all contingencies as equally credible, regardless of component age, condition, or failure history.
- It treats all loads as equally critical, regardless of the social or economic value of supply continuity.

> *In short: N-1 is a risk criterion with no probability and no consequence. It defines a boundary without describing what lies on either side of it.*

3.3 Why It Worked — and Why It No Longer Does

N-1 was developed in an era of limited computational capability, minimal real-time measurement, and relatively homogeneous, predictable power systems dominated by large synchronous generators and stable, slowly growing load. In that context, a simple, conservative, universally applicable rule produced robust and broadly acceptable results.

Today's power system is fundamentally different. High and growing shares of variable renewable generation introduce stochastic production patterns that N-1 was not designed to handle. Increased interconnection creates complex, state-dependent constraint interactions that cannot be captured by pairwise contingency analysis. Demand from data centres, AI infrastructure, and industrial electrification is growing rapidly, with connection requests that cannot be accommodated under conservative static N-1 assessments. The cost of N-1 compliance — in curtailed generation, deferred connections, and stranded investment — is now measurable and significant at a macroeconomic level (IEA, 2026).

The result is a systematic mismatch: a deterministic, context-blind criterion applied to a stochastic, highly interdependent system. In some situations, N-1 is overconservative, blocking capacity that could be used safely. In others it is insufficient, because the real risk comes from combinations of events that N-1 does not capture. The problem is not that it is too strict or too lenient. It is that it is blind to context. The Coordinated Security Assessment Methodology (CSAM) developed by ENTSO-E and approved by ACER provides a partial response, but does not resolve the fundamental probabilistic gap (ACER, 2023; ENTSO-E, 2023).

4. Quantified N-1: Seeing the System Clearly (+15–30%)

4.1 The Implicit Margin

The N-1 criterion does not eliminate risk. It displaces it into an implicit, unquantified reserve — a conservative buffer maintained in all operating states regardless of actual system conditions. This buffer exists because, without real-time probabilistic risk assessment, operators cannot distinguish between a state that is genuinely close to a security boundary and one that merely appears so under worst-case deterministic assumptions.

This implicit margin is not free. It represents foregone transfer capacity, deferred connections, and constrained dispatch. Its cost is borne by consumers through congestion, by investors through connection delays, and by society through slower electrification. International analyses consistently identify this as a primary source of latent grid capacity (IEA, 2026).

4.2 Making the Margin Visible

Probabilistic Risk Assessment (PRA) and Dynamic Security Assessment (DSA) make this margin explicit. By quantifying the actual probability-weighted consequence of contingencies — in common units such as Expected Energy Not Served (EENS) or time-at-risk — they allow operators to distinguish between:

- States where the implicit N-1 margin is genuinely needed.
- States where it is excessive, and capacity can be safely released.
- States where it is insufficient, and real risk exceeds what N-1 assumptions imply.

This is not a departure from N-1. It is a precise, quantified implementation of N-1's underlying intent: to maintain system risk within acceptable bounds. The methodology has been developed and refined in academic and operational contexts over more than two decades (Svendsen et al., 2012; Svendsen et al., 2017; Milenkovic, 2026), and its integration with dynamic stability analysis has been demonstrated to enable increased infrastructure utilisation while maintaining system security (Raab et al.).

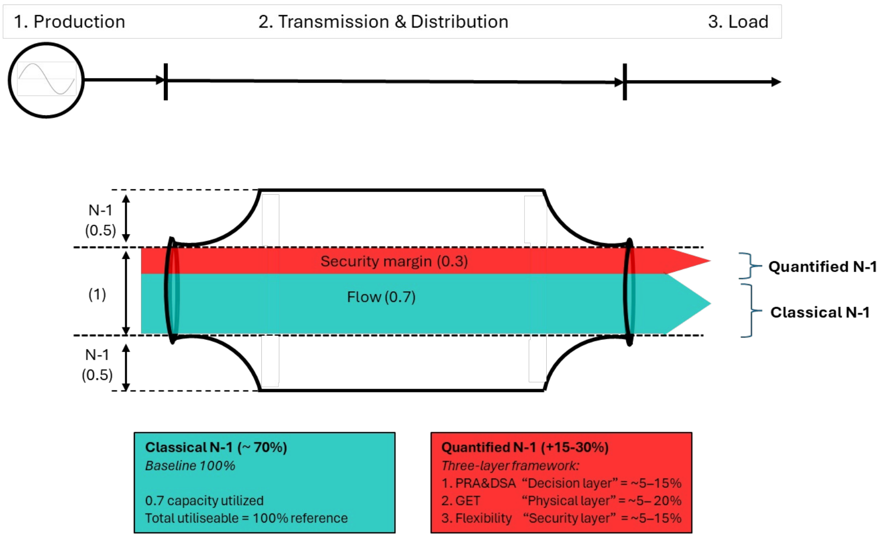

Figure 1 *Quantified N-1 increases capacity utilisation compared to classical N-1. The implicit deterministic margin (left) is made explicit and dynamic through PRA/DSA (right), releasing capacity that was previously held in invisible reserve. Conceptual illustration; values are indicative.*

4.3 The 15–30% Capacity Uplift

Based on the IEA's estimate that 1,200–1,600 GW of additional capacity could be unlocked globally through operational measures (IEA, 2026), and translating this against an approximate global utilisable transmission capacity base of 5,000–6,000 GW, the indicative system-level potential is 15–30% of current utilisable capacity (authors' interpretation).

> ***This potential does not come from new infrastructure. It comes from accessing the implicit security margin that deterministic N-1 operation holds in reserve, using quantified risk management to determine, in real time, how much of that margin can be safely released. The 15–30% range should not be interpreted as a universal forecast; it is an indicative hypothesis intended to guide system-specific validation.***

Stages 1 and 2 of the implementation pathway deliver this potential through the three-layer framework. Stage 3 — Virtual N-1 — extends capacity further through conditional load connection:

Stage	Description	Indicative Uplift
Stage 1: N-1 with open eyes	Implement PRA and DSA as the Decision Layer. Quantify the implicit N-1 security margin in real time. Release capacity where quantified risk confirms it is safe to do so. No change to the N-1 rule itself.	15–30%
Stage 2: Risk-differentiated operation	Introduce GET (Physical Layer) and flexibility mechanisms (Security Layer) coordinated against the PRA/DSA risk baseline. Extends and consolidates Stage 1 gains.	Consolidates Stage 1
Stage 3: Virtual N-1	Conditional connection for eligible load categories. Fault-event triggers initially; risk-threshold triggers as Decision Layer matures. Enlarges the effective capacity boundary.	+10–20% additional

Table 2

5. Virtual N-1: Elastic Capacity Through Conditional Connection

5.1 The Next Step

Quantified N-1 improves the utilisation of the existing N-1 security margin. It does not fundamentally change the structure of the criterion. Virtual N-1 does.

The premise is straightforward. The N-1 criterion requires that the system can withstand the loss of any single component without loss of supply to any load. This requirement is uniform: all loads, regardless of their nature, their ability to interrupt, or their economic preference for lower-cost interruptible supply, are treated identically. There is no mechanism for a load to accept a different risk profile in exchange for earlier, faster, or cheaper connection. Virtual N-1 introduces that mechanism.

5.2 Formal Definition

> ***Virtual N-1 — Formal Definition Virtual N-1 is a grid operating principle under which the transmission system is permitted to operate in an N-X state with respect to a defined category of conditionally connected load, provided that: (a) the affected load is equipped with automatic protection that disconnects or curtails it as a direct function of a quantified system risk threshold being exceeded or as an immediate response to a fault event; and (b) the resulting system state, following disconnection, satisfies the N-1 criterion for all remaining priority load. The system thereby maintains effective N-1 security through active consequence management rather than passive structural margin.***

Three elements of this definition are critical:

- Conditional connection: The load connects under a contract that explicitly accepts interruptibility as a condition of connection, not as an emergency measure.
- Automatic protection: Disconnection or curtailment is implemented through pre-installed system protection — either triggered by a quantified risk threshold (pre-fault, risk-based) or by the fault event itself (post-fault, protection-based). It is not dependent on operator action.
- Maintained N-1 for priority load: The overall system, following disconnection of conditional load, satisfies the N-1 criterion for all load that has not accepted conditional connection terms. Security for priority consumers is not reduced; it is preserved by design.

5.3 The Two Trigger Mechanisms

Virtual N-1 operates through two distinct but complementary disconnection mechanisms, which can be deployed independently or in combination:

Trigger Type	Description
Risk-threshold trigger	Conditional load is curtailed when quantified system risk — expressed as EENS or an equivalent metric — exceeds a pre-agreed threshold. This is a pre-fault, forward-looking mechanism. Typical response time: seconds to minutes via automatic control or market signal. Requires the Decision Layer (PRA/DSA) to be operational (Stage 1).
Fault-event trigger	Conditional load is disconnected as a direct, automatic response to a defined fault event — similar to conventional system protection but applied to demand rather than to network elements. Typical response time: milliseconds to seconds via protection relay. Can be implemented with existing protection technology without full Stage 1 capability.

Table 3

The distinction matters for implementation sequencing. Fault-event triggers can be deployed near-term with existing protection technology. Risk-threshold triggers require real-time PRA/DSA capability as established in Stage 1, and therefore represent a more advanced stage of implementation. Both mechanisms are analogous in principle to existing demand response and interruptible supply frameworks, such as those operated by TenneT in the Netherlands and ERCOT under Texas SB6 (2025) (Ember, 2026; Carbon Direct, 2026).

5.4 The Bottleneck — Three States

The double bottleneck illustration from White Paper 1 (Svendsen, Nyiredy & Milenkovic, 2026) — depicting the capacity constraint as two successive restrictions in a flow system — provides an intuitive representation of Virtual N-1's effect. In the conventional N-1 paradigm, both restrictions are fixed: the physical network sets one constraint, and the deterministic security margin sets another. Quantified N-1 relaxes the second restriction. Virtual N-1 goes further: by introducing conditional load as a controlled release valve, it effectively enlarges the neck of the bottleneck itself.

State	Bottleneck Description	Indicative Capacity
Classic N-1	Both restrictions are fixed: physical network + static deterministic margin. The margin is invisible and unmanageable.	100% — reference baseline
Quantified N-1 (Stages 1–2)	Physical limit unchanged. Security margin made explicit and dynamic via PRA/DSA. Second restriction becomes manageable.	115–130% of baseline
Virtual N-1 (Stage 3)	Physical limit unchanged. Security margin managed via Decision Layer. Plus: conditional load extends effective capacity beyond the existing margin. Bottleneck is structurally larger.	125–150% of baseline (conditional portion is non-firm)

Table 4

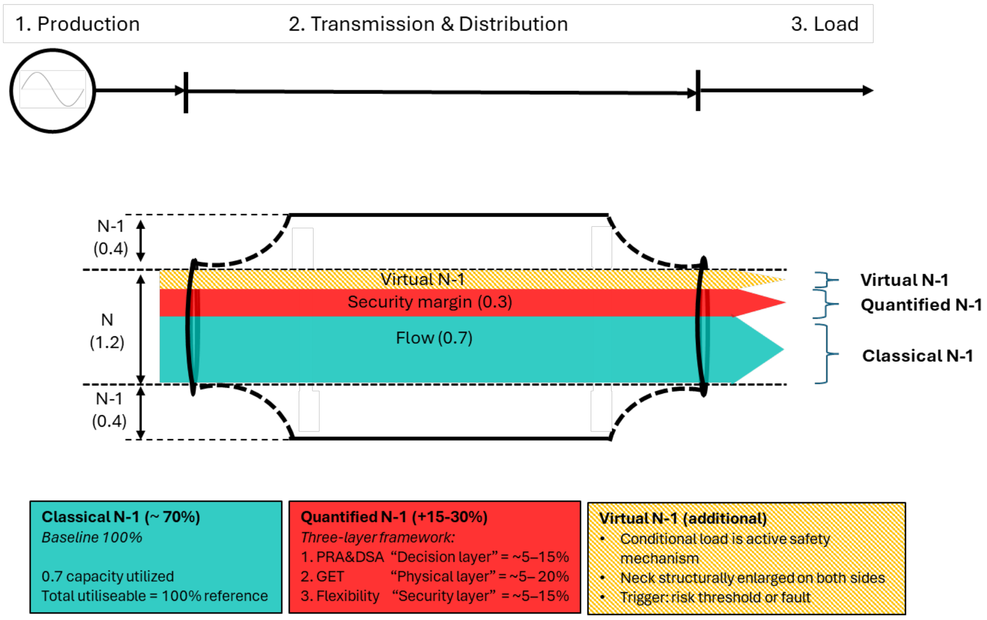

Figure 2 *Double "bottleneck" with enlarged neck through Virtual N-1. Left: classic N-1 with two fixed restrictions. Centre: quantified N-1, second restriction made manageable. Right: Virtual N-1, bottleneck neck structurally enlarged through conditional load. The Virtual N-1 section should be rendered in stippled or hatched fill to distinguish conditional from firm capacity. Conceptual illustration.*

> ***Key insight: Virtual N-1 does not reduce system security. It redistributes the consequence of potential failure — from an uncontrolled, uniform outcome affecting all consumers, to a controlled, contractually defined outcome affecting only those loads that have accepted and been compensated for conditional connection terms.***

6. Candidate Load Categories

Not all loads are candidates for Virtual N-1 connection. The operating principle is most applicable to large, flexible, interruptible loads where: (a) the economic value of connection is high enough to justify conditional terms; (b) the load can tolerate interruption without safety risk or irreversible process damage; and (c) the load volume is sufficient to make a measurable difference to system capacity.

Exclusions include critical national infrastructure (hospitals, emergency services, water treatment, communications), loads without on-site backup or UPS capability, loads with process-specific constraints that prevent rapid interruption, and residential and small commercial load.

Load Category	Characteristics Relevant to Virtual N-1
Data centres and AI infrastructure	Large, rapidly growing, seeking multi-hundred MW connections. Typically equipped with UPS and backup generation. Can implement load-shedding at sub-second timescales. Demonstrated 25% load curtailment without GPU service interruption (Carbon Direct, 2026). 5–15 GW of dispatchable grid resources identified as achievable across the US data centre fleet (Avanza Energy, 2025). Economic incentive to accept conditional terms in exchange for faster grid access. DCFlex initiative (EPRI, 2024) promotes flexible grid integration as a standard connection model (Ember, 2025).
Hydrogen electrolysers	Inherently flexible — can ramp to zero within seconds. PEM electrolysers respond faster to frequency deviations than gas turbine generators (Alaperä et al., 2020). Fully compatible with both fault-event and risk-threshold triggers. Increasing scale as hydrogen production expands under European and global policy frameworks.
Aluminium smelters and electrolytic industry	Established precedent for interruptible supply contracts in Nordic and other systems. EnPot retrofit allows 30% energy variation to match electricity supply and price fluctuations (IEA, Aluminium). Great thermal inertia makes short interruptions acceptable with manageable consequences (ScienceDirect, 2020). Fast demand response via rectifier DC voltage adjustment. Hierarchical control architectures for primary frequency support demonstrated at scale (ScienceDirect, 2020).
Battery storage co-located with large demand	Provides bridging capability during curtailment events. Can sustain host load during disconnection window while grid protection operates. Enhances commercial viability of Virtual N-1 terms for data centres and industrial facilities.

Table 5

A Lawrence Berkeley National Laboratory analysis found that 76 GW of additional load is integratable into the US system if data centres curtail for only 0.25% of maximum uptime (Knowledge Problem, 2026) — illustrating the asymmetric value of even minimal flexibility commitments.

7. Empirical Foundation: Precursors at Scale

7.1 The Connection Queue

The scale of the current connection queue globally provides the empirical starting point for the Virtual N-1 analysis. International data on queued demand is striking in both volume and composition (IEA, 2026; Ember, 2025; NESO, 2026):

Market / TSO	Queue Size	Primary Driver	Source
Italy (Terna)	30 GW demand queue by end-2024	Data centres (80% registered in last 12 months)	Ember, 2025
UK (NESO)	86 GW demand pipeline 2030–2035	Hyperscale data centres + industrial electrification	NESO, 2026
Texas (ERCOT)	40 GW projected data centre load by 2028	AI infrastructure build-out	Carbon Direct, 2026
PJM (US)	94% of 32 GW peak growth driven by data centres	AI and cloud infrastructure	PJM, 2025
Netherlands (TenneT)	Non-firm contracts unlocked 9.1 GW	Industrial + residential electrification	Ember, 2026

Table 6

The Netherlands figure is particularly significant: TenneT's non-firm contracts — a direct precursor to Virtual N-1 — already unlocked 9.1 GW equivalent to 40% of national peak electricity demand (Ember, 2026). This is not theoretical. It is demonstrated at scale. The Swedish Energy Markets Inspectorate (Ei, 2023) has reviewed conditional grid connections across multiple jurisdictions and identified a consistent pattern of both technical feasibility and commercial traction.

7.2 Existing Non-Firm and Conditional Connection Experience

Several TSOs have already implemented precursor mechanisms that provide empirical grounding for Virtual N-1 capacity estimates:

TSO / Market	Mechanism	Demonstrated Result	Relevance to Virtual N-1
TenneT (Netherlands)	Non-firm connection with curtailment obligation	9.1 GW unlocked ≈ 40% of national peak demand (Ember, 2026)	Direct analogue at national scale
NESO (UK)	GSP Technical Limits; non-firm connections; 500 MW per AI Growth Zone	Expedited connection pathway for large loads (NESO, 2026)	Structural precursor; regulatory framework buildable
PJM (US)	Non-Capacity-Backed Load (NCBL) — curtailed first in emergencies; no capacity charges	Framework for 50+ MW loads; data centres primary target (Modo Energy, 2025)	Risk-threshold trigger analogue
ERCOT / Texas (SB6, 2025)	Mandatory curtailment for loads >75 MW during grid emergencies	Legally required for large loads	Fault-event trigger at legislative level
EU (12 TSOs)	Non-firm connections; Poland, Latvia, Finland, Spain introducing frameworks	EU-wide momentum; non-firm as mainstream option (Ember, 2026)	Regulatory readiness across European markets

Table 7

8. Quantifying the Additional Capacity Potential

The incremental capacity potential from Virtual N-1 is determined by three interacting factors. Each introduces a range of uncertainty; the final estimate is the product of all three. The approach is consistent with the methodology applied to derive the 15–30% estimate in White Paper 1 (Svendsen, Nyiredy & Milenkovic, 2026) — an indicative hypothesis grounded in empirical evidence, not a modelled forecast.

Factor A — Eligible Load Share

What proportion of new load seeking connection is technically and commercially suitable for Virtual N-1 terms?

Scenario	Eligible Share of New Large-Load Queue	Basis
Conservative	30–40%	Only large hyperscale data centres (>100 MW) and established interruptible industrial loads. Excludes most data centres below 50 MW and newer load categories.
Moderate	45–55%	Hyperscale + mid-size data centres willing to accept conditional terms for faster connection + electrolysers + established aluminium/industrial.
Optimistic	55–65%	Above plus AI-native facilities designed for flexibility from the outset + BESS co-located loads. Consistent with DCFlex ambitions and EPRI/industry direction (Ember, 2025).

Table 8

Central estimate for Factor A: 40–55% of new large-load queue is Virtual N-1-eligible under realistic commercial conditions.

Factor B — Committed Curtailment Depth

What percentage of connected load can the Virtual N-1 consumer credibly commit to curtailing, at the response speed required by the trigger mechanism? The committed curtailment must be sufficient to restore N-1 compliance for priority load following the trigger event.

Curtailment Tier	Depth	Response Time	Load Categories	Trigger Compatibility
Tier 1 — Deep fast	20–30%	Milliseconds (protection relay)	Aluminium smelters, electrolysers (ScienceDirect, 2020; Alaperä et al., 2020)	Fault-event trigger. No DSA required.
Tier 2 — Moderate fast	15–25%	Seconds (automatic control)	Hyperscale data centres with UPS + backup generation (Carbon Direct, 2026)	Both triggers.
Tier 3 — Moderate managed	10–20%	Minutes (automated market signal)	Mid-size data centres, batch-compute facilities (ScienceDirect, 2025)	Risk-threshold trigger only.
Tier 4 — Shallow	5–10%	Minutes to hours	Industrial loads with limited flexibility	Risk-threshold only. Limited system value.

Table 9

Central estimate for Factor B: a weighted average committed curtailment depth of 15–22% across the eligible load mix, with Tier 1 and Tier 2 loads providing the majority of system value.

Factor C — System Headroom Translation

How does a given committed curtailment depth translate into additional connected load, expressed as a percentage of utilisable grid capacity? In a typical meshed transmission system, N-1 compliance requires maintaining approximately 15–25% headroom on critical circuits. A Virtual N-1 load with committed 20% curtailment can be connected to a circuit operating at 85% of N-1 limit, because a curtailment event reduces circuit loading by 20% — restoring compliance. The load “pays” for its own N-1 margin through the curtailment commitment, rather than the system operator holding that margin in reserve.

> *Illustrative moderate scenario: Queued load as % of current capacity: 30% × Eligible share: 50% × Translation efficiency: 80% = 12% incremental capacity uplift. (Authors’ interpretation.)*

Scenario Synthesis

Parameter	Conservative	Moderate	Optimistic
Queued load as % of current utilisable capacity	20%	30%	40%
Eligible share (Factor A)	35%	50%	60%
Average committed curtailment (Factor B)	12%	18%	25%
Required N-1 margin (system-dependent)	20%	18%	15%
Incremental Virtual N-1 uplift	~5%	~10–15%	~20–25%

Table 10

> ***Recommended indicative range: +10–20% of utilisable grid capacity, beyond the 15–30% from the three-layer framework. Empirical grounding: TenneT’s demonstrated 9.1 GW ≈ 40% of national peak demand from non-firm contracts (Ember, 2026). This range should be treated as an indicative hypothesis requiring system-specific validation.***

9. The Cumulative Capacity Picture

The three-layer framework and Virtual N-1 are complementary but not fully independent. They share a common prerequisite — the Decision Layer (PRA/DSA) — and some overlap exists at the margin. The conservative cumulative estimate accounts for this overlap (authors' interpretation, derived from IEA, 2026):

Layer	Mechanism	Indicative Uplift	Cumulative (Conservative)
Baseline	Deterministic N-1 operation today	—	0% (reference)
Three-layer framework (Stages 1–2)	PRA/DSA + GET + Flexibility Makes implicit margin explicit and manageable	+15–30%	+15–30%
Virtual N-1 (Stage 3)	Conditional load connection Load serves as active safety mechanism	+10–20%	+25–50%
Combined (with overlap adjustment)	Full three-stage implementation	+20–45%	+20–45%

Table 11

The overlap adjustment (subtracting 5–10 percentage points from simple addition) reflects the fact that some of the load that Virtual N-1 enables would partially overlap with the capacity released by the three-layer framework under optimistic assumptions. Under conservative assumptions, the overlap is minimal and the potentials are largely additive.

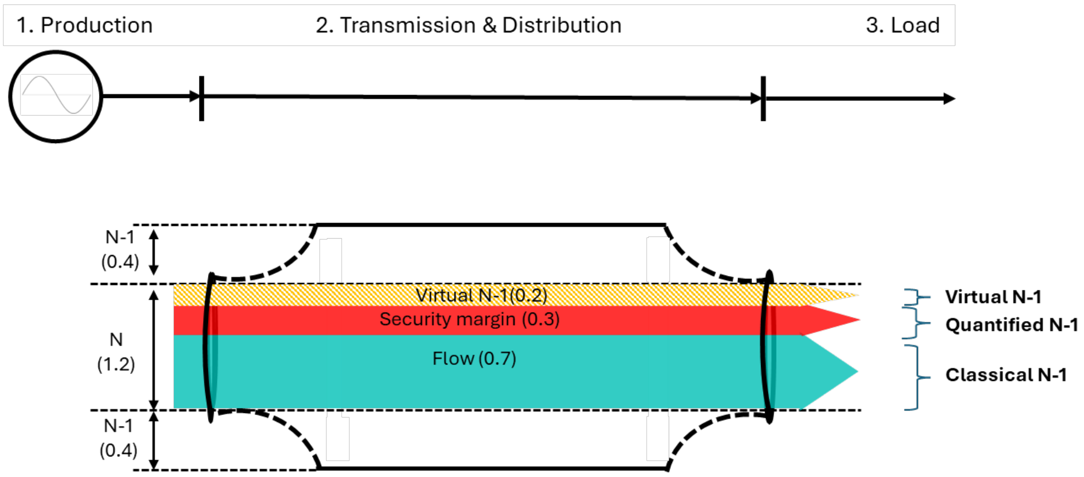

Classical N-1 (~ 70%)
Baseline 100%

0.7 capacity utilized
Total utiliseable = 100% reference

Quantified N-1 (+15-30%)
Three-layer framework:
1. PRA&DSA "Decision layer" = ~5–15%
2. GET "Physical layer" = ~5– 20%
3. Flexibility "Security layer" = ~5–15%

Virtual N-1 (+10-20% additional)
- Conditional load is active safety mechanism
- Neck structurally enlarged on both sides
- Trigger: risk threshold or fault

Cumulative potential

Baseline:	100%
+Three-layer:	+15-30%
+Virtual N- 1:	+10-20%
Total:	+25-50%

Total realistic system level potential (non-additive, overlapping): 15-30% (three-layer) + 10-20% (Virtual N-1) = 20-45%
Accounting for overlap of released load from three-layer framework and Virtual N-1.

Figure 3 *Cumulative capacity potential — three states stacked. Baseline (deterministic N-1) / Three-layer framework (+15–30%) / Virtual N-1 (+10–20% additional). Combined indicative range: +20–45% of utilisable capacity. Virtual N-1 section shown in stippled fill to indicate conditional (non-firm) nature. Conceptual illustration; values* ***are indicative and should not be interpreted as additive or independently realisable.***

Note that the Virtual N-1 portion of capacity is non-firm: it is real capacity that can be reduced if a trigger event occurs. This distinction must be preserved in system planning, reliability indices (LOLP, EENS), and investment economics. Virtual N-1 load cannot be relied upon as firm demand for generation adequacy assessments without specific contractual and technical assurances.

10. Regulatory and Liability Framework

10.1 From Uniform to Differentiated Security

Current regulatory frameworks treat all connected load as entitled to the same level of supply security, defined by N-1 compliance. Virtual N-1 requires a regulatory distinction between priority load (subject to full N-1 security, as today) and conditional load (connected under explicit contractual terms specifying interruptibility conditions, trigger mechanisms, and compensation arrangements).

This distinction is not new in principle — interruptible supply contracts have existed in many systems for decades, and the Swedish Energy Markets Inspectorate has reviewed conditional connection frameworks across multiple jurisdictions (Ei, 2023). What is new is the systematic application of this principle as a grid planning and connection tool, rather than a demand response measure of last resort.

10.2 Liability and Contractual Framework

The contractual framework must address:

- Who bears the cost of a disconnection event, and under what conditions.
- What compensation the conditional load receives for accepting interruptibility terms.
- What obligations the TSO has to minimise the frequency and duration of disconnection events.
- How the risk-threshold trigger level is set, reviewed, and audited.

Establishing this framework is a regulatory and commercial task, not a technical one. The technical mechanisms for implementing Virtual N-1 exist. The institutional infrastructure for governing it does not yet, in most jurisdictions. Existing precedents — PJM's Non-Capacity-Backed Load framework, ERCOT's SB6 mandatory curtailment requirements, and TenneT's non-firm agreements — provide practical starting points (Modo Energy, 2025; Carbon Direct, 2026; Ember, 2026).

10.3 European Regulatory Pathway

Body	Required Action
ACER and the European Commission	Recognition of conditional connection categories in the network codes; clarification of how Virtual N-1 interacts with SOGL and CACM frameworks (ACER, 2023; European Commission, 2025).
ENTSO-E	Development of a common methodology for quantifying risk thresholds used in trigger mechanisms; coordination across TSO boundaries for cross-border conditional connections. Building on existing CSAM work (ENTSO-E, 2023).
National regulatory authorities	Approval of conditional connection tariffs; definition of eligible load categories and interruptibility terms; liability framework for disconnection events. Netherlands and UK provide regulatory templates.

Table 12

10.4 Implementation Timeline

Horizon	Milestone
Near-term (1–3 years)	Regulatory frameworks for conditional connection in early-adopter jurisdictions (Netherlands, UK). Fault-event trigger pilots with data centres and industrial loads. No Decision Layer required.
Medium-term (3–5 years)	Meaningful Virtual N-1 capacity contribution from fault-event trigger deployments. Decision Layer (Stage 1) operational in several TSOs, enabling risk-threshold trigger design.
Full implementation (5–10 years)	Risk-threshold trigger deployment at scale, requiring Stage 1 maturity. Cross-border conditional connection coordination under ENTSO-E methodology.

Table 13

11. Caveats and Qualifying Conditions

The estimates in this paper are indicative and subject to the following qualifying conditions, which must be stated explicitly wherever the figures are cited.

11.1 System Specificity

The achievable Virtual N-1 uplift varies significantly across systems. Key system-specific factors include network topology (radial systems offer less headroom than meshed systems), composition of connection queue, existing N-1 margin utilisation, and the maturity of the regulatory framework. Systems with established non-firm connection frameworks — Netherlands, UK — can implement Virtual N-1 faster than those starting from scratch (Ei, 2023).

11.2 The Conditional Nature of Virtual N-1 Capacity

Virtual N-1 capacity is not equivalent to firm capacity. The distinction matters for planning (generation adequacy assessments), investment (revenue streams subject to curtailment risk), and reliability metrics (LOLP, EENS must account for the conditional nature of the load). These caveats are manageable within the framework proposed in this paper, but they require explicit treatment in regulatory accounting and system planning.

11.3 Non-Additivity with the Three-Layer Framework

The Virtual N-1 uplift and the three-layer framework uplift are not fully additive. The Decision Layer (PRA/DSA) is a shared prerequisite for both. Systems that have not yet implemented Stage 1 cannot implement Stage 3 with risk-threshold triggers. Fault-event triggers can be implemented earlier and without the full Decision Layer, providing a partial Virtual N-1 benefit at lower system maturity.

11.4 Methodological Status

> *The +10–20% incremental Virtual N-1 estimate is an indicative hypothesis, derived from a structured analysis of empirical evidence and physical system constraints. It is not a modelled result, a universal forecast, or a transferable capacity figure. The IEA's 1,200–1,600 GW global estimate (IEA, 2026) covers a broad range of operational measures and does not separately quantify a Virtual N-1 contribution. The estimates in this paper are developed independently using bottom-up reasoning from empirical load data and system physics, and are broadly consistent with but not derived from the IEA figures. These figures should be used as a starting point for system-specific analysis, not as results to be cited without qualification.*

12. Conclusions and Specific Recommendations

The global electricity grid is constrained not primarily by physics or technology, but by the operating criteria under which it is managed. N-1 — a practical heuristic developed in the 1950s for a simpler and more predictable system — has become the binding constraint on grid utilisation in an era of variable renewables, exploding electrification demand, and rapid load growth (IEA, 2026).

This paper has traced the argument from first principles: from what a power system actually does (manage risk under uncertainty), through what N-1 actually is (a deterministic proxy for risk, not a risk measure), to how quantified risk management can unlock 15–30% more utilisation within the existing N-1 paradigm, and finally to Virtual N-1 — a new operating principle that extends that potential by a further 10–20%.

> ***The central claim: Virtual N-1 does not weaken power system security. It makes security explicit, differentiated, and manageable — replacing an implicit, uniform, and increasingly expensive buffer with a transparent, contractually governed, and economically efficient alternative.***

Specific Recommendations

- For TSOs: Begin mapping eligible load share in existing connection queues. Identify data centre, electrolyser, and industrial load with credible curtailment capability. Initiate pilot designs for fault-event trigger mechanisms with early-mover customers. Engage with ENTSO-E on cross-border coordination methodology.
- For National Regulatory Authorities: Develop conditional connection tariff frameworks defining eligible load categories, interruptibility terms, and compensation structures. Use Netherlands and UK frameworks as regulatory templates. Engage with ACER on network code implications (European Commission, 2025).
- For ACER and ENTSO-E: Prioritise development of a common methodology for risk-threshold trigger quantification, building on CSAM (ENTSO-E, 2023; ACER, 2023). Define how Virtual N-1 conditional load categories interact with cross-border capacity allocation and security coordination frameworks.
- For large load customers (data centres, electrolysers, industrial): Assess commercial value of Virtual N-1 connection terms relative to firm connection queue timelines. Engage with TSOs on pilot participation. Design on-site protection infrastructure compatible with fault-event trigger requirements. Consider DCFlex-aligned connection models (Ember, 2025).
- For policymakers: Recognise that the loads seeking Virtual N-1 connection are already in the queue — the regulatory frameworks needed to connect them do not yet exist in most jurisdictions. Accelerating that institutional infrastructure is as important as building new grid hardware (IEA, 2026).

12.1 Key Takeaway: We Need a Bigger House

The electricity system is running out of room. Connection queues are growing, generation is waiting, and new demand — from data centres, electrolysers, and industrial electrification — is arriving faster than the grid can accommodate it. In simple terms: we need more capacity real estate.

The instinct is to build new infrastructure. But before we build new, we should ask whether the house we already have is being used to its full potential. In most power systems today, it is not.

Think of the grid's available capacity as a house. Classic N-1 operation gives us a single-storey building — functional, safe, but low. The ceiling is fixed by a rule written decades ago, and nobody questions how much headroom is actually needed above it. The margin is there, invisible and unmanaged, holding the ceiling lower than it needs to be.

To build higher, we need a foundation we can trust. That is what the Decision Layer does — PRA and DSA give us real-time, quantified knowledge of how the structure is actually performing, not just a conservative estimate of how bad things might get. With a solid foundation in place, we can safely add floors. The first floor — Grid Enhancing Technologies — makes better use of what the physical structure can carry. The second floor — flexibility and markets — manages the residual load so the structure is never overloaded. Together, these three layers give us a taller, more efficient building. And crucially, we build this under existing regulations. No new planning permission required. The N-1 criterion remains intact. We are simply using the building more intelligently.

This is what Stages 1 and 2 deliver: a 15–30% increase in utilisable capacity, within current rules.

But there is a ceiling. And at some point, the question shifts from how do we use the existing building better to how do we get permission to build the final floor.

Virtual N-1 is that permission — granted not by relaxing the building code, but by introducing a new category of occupancy. Some tenants, in exchange for lower rent and faster access, agree to vacate on short notice if the structure needs to shed load. They sign a contract. They install the right equipment. And because their departure is orderly, automatic, and bounded, the building can safely accommodate them on a floor that would otherwise have remained empty — the loft, with its own conditional ceiling.

The building code — N-1 — is not broken. Security for all standard occupants is fully preserved. What changes is that the regulator recognises a new kind of tenant, and a new kind of floor.

The result: a building that is 20–45% taller than the one we started with. Same plot. Same foundations. Smarter use of the space — and one carefully governed extension at the top.

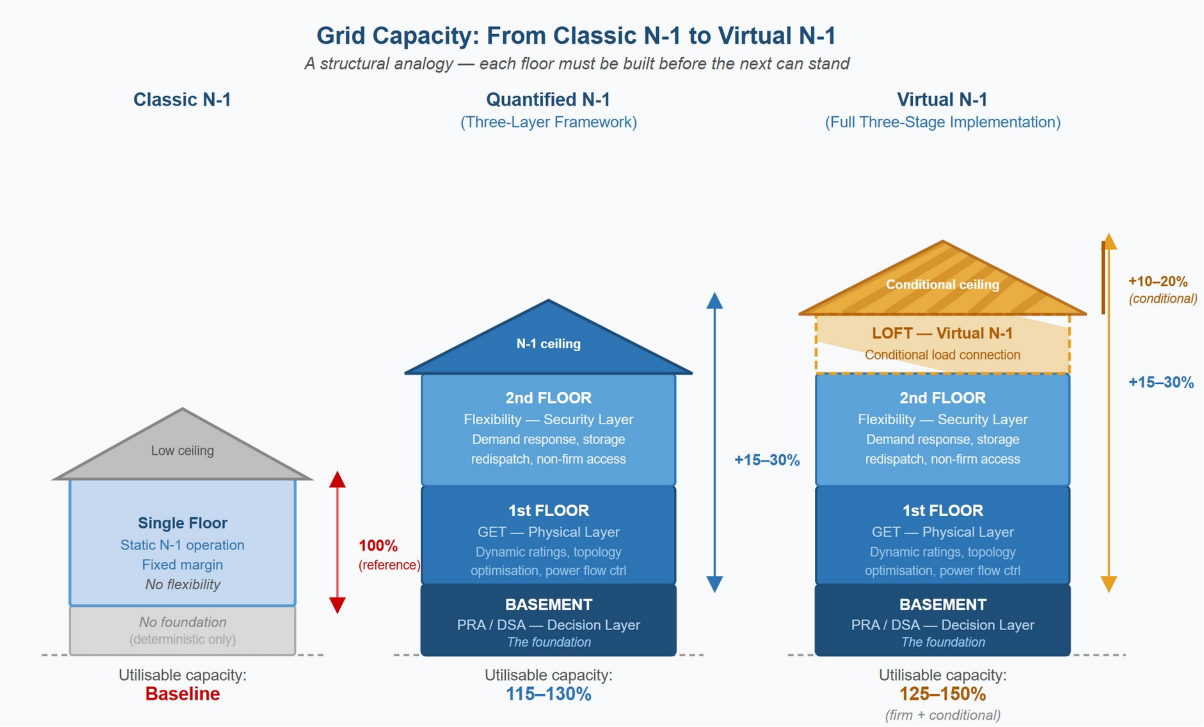

Figure 2 *Grid Capacity: From Classical N-1 to Virtual N-1 - A structural analogy - each floor must be built before the next*

References

ACER (2023). Decision on the Coordinated Security Assessment Methodology (CSAM). Agency for the Cooperation of Energy Regulators, Ljubljana.

Alaperä et al. (2020). Grid Balancing with a Large-Scale Electrolyser Providing Primary Reserve. ResearchGate.

Avanza Energy (2025). Load Flexibility’s Breakthrough Moment — Driven by Data Centers.

Carbon Direct (2026, March). The billion-dollar case for enabling data center load flexibility.

Ei — Swedish Energy Markets Inspectorate (2023). Conditional grid connections — a literature review. EiR2023:08.

Ember (2025, September). Grids for data centres: ambitious grid planning can win Europe’s AI race.

Ember (2026, April). Crossed wires: Grid capacity could block EU energy security.

ENTSO-E (2023). Coordinated Security Assessment Methodology (CSAM). Brussels.

ENTSO-E (2023). Biennial Progress Report on Probabilistic Security Assessment and Risk Management. Brussels.

European Commission (2025). European Grids Package (COM/2025/1005). Brussels.

IEA (2026). Electricity 2026. International Energy Agency, Paris.

IEA (2023). Electricity Grids and Secure Energy Transitions. International Energy Agency, Paris.

IEA. Aluminium — Energy System. iea.org/energy-system/industry/aluminium.

Knowledge Problem (2026). Data Centers, Flexibility, and the Architecture of the Grid.

Milenkovic, M. (2026). Application of Markov Decision Processes for risk-informed power system operation (unpublished internal work).

Modo Energy (2025). PJM’s answer to large load: Curtailment as a market lever.

NESO (2026). NESO reveals results of UK’s grid connection reforms. Data Centre Dynamics.

Raab, A., et al. (n.d.). Security assessment for higher loaded power system operation to 2030. Siemens AG / TenneT / FAU.

Rausand, M. and Høyland, A. (2004). System Reliability Theory: Models, Statistical Methods, and Applications. John Wiley & Sons.

ScienceDirect (2020). Hierarchical control of aluminum smelter loads for primary frequency support.

ScienceDirect (2025). Data centres as a source of flexibility for power systems.

Svendsen, A.B. (2002). Reliability Analysis for Protection and Control Systems in Power Systems. Master Thesis, NTNU, Norway.

Svendsen, A.B. et al. (2012). Online Reliability Assessment of Power Systems. PMAPS Conference.

Svendsen, A. B., Tollefsen, T., Pedersen, R. F., Petursson, K. P., Patel, D., & Lampe, O. D. (2014). Abstract representation of power system networks as a function of regularity properties. ESREL 2014.

Svendsen, A.B. et al. (2017). Modelling Weather Dependence in Online Reliability Assessment of Power Systems. Journal of Risk and Reliability.

Svendsen, Nyiredy & Milenkovic (2026). Unlocking Grid Capacity Through Risk-Based Operation. White Paper 1.

From N-1 to Virtual N-1 · White Paper 2 · Svendsen, Nyiredy**, Milenkovic** · May 2026*