SLTA | SWISS LEGALTECH ASSOCIATION

LAW&TECH INSIGHTS MAGAZINE

2026-10

SLTA TECH & LAW INSIGHTS

THE UNKNOWN PERIMETER

OPERATIONAL OPACITY AND CYBER RESILIENCE OF CRITICAL INFRASTRUCTURE: THE ELECTRICITY INDUSTRY

SANDRINE VIRGINIE HILAIRE, DIGITAL GOVERNANCE AND TRANSFORMATION
AND
MOHAMMAD USMAN, INTERNATIONAL RELATIONS AND DIPLOMACY

Editorial Team:
Rucsandra Bratu and Ana Maraqvelidze

www.swisslegaltech.ch

INTRODUCTION

Over the last four years, European cybersecurity law has transformed what was previously a technical expectation into a legal requirement. Directive (EU) 2022/2555 ("**NIS 2**") requires essential and important entities to adopt "asset management" measures[1] in relation to cybersecurity risk management. The sector-specific network code on cybersecurity for cross-border electricity flows ("**NCCS**", Commission Delegated Regulation (EU) 2024/1366) obliges designated entities to perform risk management for all assets within a defined high-impact or critical-impact perimeter.[2] Last but not least, the Cyber Resilience Act ("**CRA**", Regulation (EU) 2024/2847) mandates that manufacturers identify and document the components of products with digital elements, including through a software bill of materials.[3] All three instruments assume that a regulated entity is able to identify what it operates.

Notwithstanding that assumption, complex organisations may fail to produce reliable knowledge of their own operational architecture for reasons unrelated to negligence, under-investment or immature tooling. Ordinary declarative methods are constrained by the organisational structures of the entity itself. Specialisation, decentralisation, delegated technical decisions and the local resolution of problems all enable organisations to operate at scale, while making their architecture increasingly difficult to observe as a whole. A compliance regime based on self-declaration of a perimeter that the declaring entity cannot fully observe may therefore produce formally compliant but substantively unverifiable documentation.

The electricity industry delivers the sharpest illustration. It combines the densest regulatory overlay in the European Union with the widest gap between declared and operational architecture for reasons that are specific to Operational Technology ("**OT**") environments rather than general to enterprise information technology. Part I of this article sets out the regulatory presupposition of self-knowledge. Part II explains why that presupposition fails structurally. Part III identifies what is specific to OT and to the electricity industry. Part IV examines the time asymmetry created by the Cyber Resilience Act in relation to the installed base. Part V assesses the legal consequences. Part VI considers the consequences of the same problem at the level of the State, where the operator is the first layer through which national infrastructure is observed.

1. THE REGULATORY PRESUPPOSITION OF SELF-KNOWLEDGE

a. NIS 2: asset management as a legal obligation

Article 21(1) of NIS 2 requires essential and important entities to take appropriate and proportionate technical, operational and organisational measures to manage the risks posed to the security of their network and information systems. Article 21(2) then specifies a minimum content of ten items, the ninth of which is "human resources security, access control policies and asset management".[4] The obligation is outcome-based rather than prescriptive. The Directive states what shall be achieved rather than how.

Two features of the architecture of NIS 2 make the asset management requirement more consequential than its brief drafting suggests. First, Article 20(1) requires management bodies to approve the risk-management measures and to oversee their implementation. It further provides that those bodies may be held liable for infringements of Article 21.[5] Cybersecurity knowledge is thereby converted into a matter of directors' duties. Second, when assessing supply chain measures under Article 21(2)(d), entities shall take into account the vulnerabilities specific to each direct supplier and service provider together with the general quality of their products and cybersecurity practices.[6] Such an assessment presupposes the identification of all relevant suppliers.

b. The Network Code on Cybersecurity: the perimeter as a legal object

Adopted in March 2024 and in force since June 2024, the NCCS goes considerably further than NIS 2 in specificity.[7] It applies only to entities designated by national competent authorities as high-impact or critical-impact based on an Electricity Cybersecurity Impact Index ("**ECII**").[8] Once designated, each entity shall perform cybersecurity risk management for all assets within its high-impact and critical-impact perimeters. That exercise comprises context establishment, risk assessment, risk treatment and risk acceptance. It shall be repeated every three years.[9] Within twelve months of designation and every three years thereafter, the entity shall report to the competent authority the mitigation controls selected, an estimate of residual risk and a list of the critical ICT service providers involved in its critical-impact processes.[10] For critical-impact entities, compliance shall be demonstrated through a national verification scheme or an independent third-party audit, with the full verification scope covered at least once every three years.[11] Supply chain controls apply across the entire lifecycle of ICT products, services and processes.[12]

The novelty of the network code lies in the perimeter. It does not merely require security measures. It requires the entity to (a) draw a boundary around a subset of its own estate, (b) declare that boundary to a regulator, and (c) accept audit against it. The perimeter is a legal object constructed by the regulated entity itself. Its accuracy is a condition of the validity of everything built upon it: the risk assessment, the control selection, the audit scope and the residual risk statement submitted to the authority.

c. The Cyber Resilience Act and the limits of the bill of materials

The CRA approaches the same problem from the product side. Annex I, Part II, point (1) requires manufacturers of products with digital elements to identify and document the components contained in the product, including by drawing up a software bill of materials ("**SBOM**") in a commonly used, machine-readable format covering at least the top-level dependencies of the product.[13] Reporting obligations for actively exploited vulnerabilities apply from 11 September 2026. The essential requirements, including the SBOM obligation, apply from 11 December 2027.[14]

The SBOM is a genuine advancement, whose limits are worth stating precisely. An SBOM describes the composition of a product as shipped. It says nothing about how that product has been deployed, what it has been connected to, which of its interfaces have been enabled, or what undocumented integration has grown around it in the fifteen years since commissioning. The CRA improves the legibility of the components, whereas the opacity that concerns this article is located in the space between them.

2. WHY THE PRESUPPOSITION FAILS

a. Declared architecture and operational architecture

Within any large organisation, several architectures exist and diverge: the documented architecture, the deployed architecture, the runtime architecture, the perceived architecture, the governance architecture, the emergency or fallback architecture, and the architecture that remains after several years of patches and workarounds. Architecture repositories become unreliable not because they are badly maintained but because they record declarative truth rather than operational truth.

A specific illustration may assist. An official diagram may show a customer relationship management system connecting to an enterprise resource planning system through an API gateway. The operational reality may be a spreadsheet export, a script running on an employee's laptop, an unmanaged software-as-a-service connector, a shared mailbox, and a batch import. Each link in that chain was created by an identifiable person, at an identifiable moment, in order to solve a real problem. None of it appears in any diagram. All of it is load-bearing. No discovery engine easily reconstructs the whole chain because part of the chain exists outside formal infrastructure altogether.

b. A structural rather than an accidental failure

The organisational cause is well established in literature. Lawrence and Lorsch demonstrated that the more uncertain and complex an organisation's environment, the more it must differentiate internally into specialised sub-units and the more it must therefore invest in integration methods to prevent fragmentation.[15] Within contemporary infrastructure operators, that differentiation has been pushed to an extreme. Integration has become structurally difficult not for want of will but for want of visibility over what actually exists.

The resulting pathology is the silo. Consider a unit organised around a technology that was, when adopted, the most appropriate response to a need. That unit will orient its efforts towards optimising the technology even as the technology becomes obsolescent, because the unit's identity and competence are inseparable from what it was built to maintain. The opacity that results is not chaos. It is the accumulated rational behaviour of units, each optimising its own segment of a system that nobody sees whole. Michael Porter's distinction between primary and support activities helps explain why this escapes executive attention. Transversal functions such as infrastructure, technology development and legal affairs traverse the whole chain without belonging to any single stage of it. They are for that reason systematically under-represented in strategic analysis.[16]

c. Shadow integration, single points of failure and the bus factor

Two operational concepts deserve legal attention. A single point of failure is an element so central that its loss stops everything depending upon it. The most consequential single points of failure in mature organisations are not technical, but rather human: the one person who understands how a critical script works or who manually operates the mailbox that triggers a scheduled import. A related notion is the bus factor, meaning the number of people whose immediate absence would collapse a critical function. It is frequently one and almost always invisible to management. It is produced by informal problem-solving in which someone resolves a difficulty outside governance. The solution becomes critical, and nobody else ever learns how it works.

The World Economic Forum's Cyber Resilience Compass reflects the same finding from the practitioner side. It identifies as a front-line practice that chief information security officers working with business units should "identify critical dependencies and single points of failure by improving visibility and developing an accurate mapping". It notes further that such mapping must recognise that supply chains are non-linear and that issues can propagate both upstream and downstream.[17] This identification as a front-line practice rather than as a baseline indicates that it is not yet general.

3. WHAT IS SPECIFIC TO OT AND TO ELECTRICITY

Three features distinguish the electricity industry from general enterprise IT and make the epistemic gap materially wider.

The first is the direction of change. Electricity systems are moving from centralised generation towards distributed complexity in the form of smart grids, distributed energy resources and prosumers. They are moving at the same time from static infrastructure towards software-defined systems with digital twins, industrial internet of things devices and AI-powered analytics layered onto plant commissioned decades earlier. Interconnection across IT and OT environments has increased together with the number of paths by which a local compromise becomes a cascading disruption. Demand growth compounds the stakes. The International Energy Agency forecasts European Union electricity demand growing by around 2% per year to 2030 and United States demand by nearly 2% annually, roughly half of the American increase being attributable to data center expansion.[18] Electricity is the critical infrastructure upon which all other critical infrastructures depend.

The second is the constraint of the environment itself. Energy systems are safety-critical, time-sensitive and constrained by legacy hardware, long equipment lifecycles and strict availability requirements. Digitalised energy systems therefore differ from conventional IT. Security solutions must be designed with explicit regard to operational requirements without risking unintended operational disruption.[19] This carries a direct epistemic consequence. In OT environments, much of the tooling that produces asset knowledge in ordinary IT is either prohibited or operationally unacceptable, including active scanning, agent deployment or credentialed enumeration. Knowledge of the estate must frequently be inferred from passive observation rather than obtained by interrogation. The organisation is not merely under-informed. It is constrained by the methods through which it may become informed.

The third is the empirical record. In December 2015, an intrusion into three Ukrainian regional distribution companies interrupted supply to approximately 225,000 customers through the remote operation of breakers at some thirty substations. It demonstrated that OT consequences follow from IT-side compromise.[20] The May 2023 campaign against the Danish energy sector is more instructive for present purposes. Attackers exploited a firewall vulnerability across 22 organisations in a coordinated first wave, followed by a second wave using vulnerabilities that were unpatched at the time. Several operators were obliged to move to island mode.[21] What made sector-wide detection possible was a shared sensor network operated by the sectoral computer emergency response team. Detection came from an external observation layer rather than from the internal architectural documentation of the affected entities.

4. THE TEMPORAL ASYMMETRY: THE CYBER RESILIENCE ACT AND THE INSTALLED BASE

The CRA is the instrument most frequently invoked as the answer to component-level opacity. It deserves closer attention than that shorthand allows because in the electricity industry its transitional architecture produces an asymmetry that runs directly against the thesis of this article.

a. Grandfathering and the two-speed estate

Article 69(2) of the CRA provides that products with digital elements placed on the market before 11 December 2027 are subject to the Regulation only if, from that date, they are subject to a substantial modification.[22] The carve-out is not total. Article 69(3) disapplies the grandfathering in respect of the reporting obligations under Article 14 which extend to products placed on the market before that date and take effect from 11 September 2026.[23] The essential requirements attach prospectively to what is placed on the market from December 2027 onwards. Those requirements include the obligation to identify and document components and to draw up a software bill of materials.

Set that against the physical reality of an electricity estate. Energy systems are constrained by legacy hardware and long equipment lifecycles. Security solutions must accommodate operational requirements that conventional IT does not face.[24] The International Energy Agency records that power systems encounter "rising risks from ageing infrastructure, extreme weather events, cyberthreats and other emerging vulnerabilities".[25] Substations, protection relays, remote terminal units and control systems commissioned in the 1990s and 2000s will still be carrying load in the 2040s. The CRA will therefore produce, within a single operator, a two-speed estate: a documented SBOM-bearing layer of recently procured equipment above an undocumented substrate. That substrate will remain quantitatively dominant for at least a further decade. In the case of the longest-lived plant, it will remain so considerably longer. The Regulation thus improves the legibility of the portion of the estate that was already best documented but without reaching the portion in which the opacity resides.

b. Support periods and the closing of the documented window

A second temporal effect compounds the first. Article 13(8) requires manufacturers to determine a support period indicating the length of time the product is expected to be in use. It shall not be less than five years unless the anticipated lifetime is shorter.[26] Five years is a floor rather than a norm. The provision expressly directs manufacturers to consider the nature and intended purpose of the product. For grid equipment, the resulting periods should be considerably longer. The structural point holds, though. The support period is calibrated to the manufacturer's commercial horizon, while the asset is calibrated to the physical life of a substation. Where the two diverge, the component does not disappear from the estate at the end of support. It reverts to the substrate. Its SBOM becomes an archival record describing a configuration that has since been patched, integrated and modified in ways the manufacturer no longer tracks.

c. The interaction with the network code

The two European instruments most relevant to the electricity industry do not treat legacy consistently. Under Article 26(2) of the NCCS, a designated entity shall perform cybersecurity risk management for all assets within its high-impact and critical-impact perimeters, without distinction as to when those assets were placed on the market.[27] By contrast, the CRA exempts pre-2027 units from its essential requirements. The operator is therefore required by the network code to assess and treat risk across an estate whose older portion was deliberately left unlit by the regulation. Compliance with the CRA will not discharge the NCCS obligation in respect of that portion. It will simply not speak to it.

It should be said in fairness to the drafters that the alternative was worse. Retroactive application to the installed base would have been unworkable. The *"substantial modification"* trigger is drafted so that standard security patching and bug fixing do not bring a legacy product into scope, precisely in order not to deter the maintenance that keeps old equipment safe. The objection advanced here is directed at an inference sometimes drawn from that regime rather than at the regime itself, namely that the CRA resolves the asset knowledge problem. What it improves is prospective, situated at the component level and on the manufacturer's side of the boundary, whereas the problem examined in this article is retrospective, situated at the integration level and on the operator's side. Regulation can require that new products arrive documented. It cannot legislate knowledge of what is already installed. If such knowledge is to exist at all, it must be reconstructed. Reconstruction from traces is a different discipline from declaration.

5. LEGAL CONSEQUENCES OF AN UNVERIFIABLE PERIMETER

If the foregoing is accepted, four consequences follow for the practitioner.

First, the reasonableness standard becomes indeterminate. Article 21(1) of NIS 2 requires measures "appropriate and proportionate" to the risks posed. Proportionality is assessed against the entity's exposure to risk, its size and the likelihood and severity of incidents. Each of those variables is a function of the estate. An entity that has under-scoped its perimeter will have calibrated its measures against an understated risk profile while remaining formally compliant on the face of its documentation. The defect is not detectable within the four corners of the compliance file.

Second, governance liability attaches to knowledge that the management body cannot independently verify. Article 20(1) requires management bodies to approve risk-management measures and provides that they may be held liable for infringements of Article 21.[28] In respect of essential entities, Article 32(5) permits competent authorities to suspend an authorisation and to impose a temporary prohibition on the exercise of managerial functions. Article 34 sets maximum administrative fines of at least €10,000,000 or 2% of total worldwide annual turnover, whichever is higher.[29] A director approving an asset management policy is approving a representation about a system that no individual within the organisation observes in full. The prudent response is not to decline the approval, but rather to document the basis of the representation together with its acknowledged limits.

Third, the incident reporting timetable presupposes architectural knowledge that the incident itself disproves. Article 23 of NIS 2 requires (a) an early warning within 24 hours of awareness of a significant incident, (b) a fuller notification within 72 hours including an initial assessment of severity and impact, and (c) a final report within one month. Under the NCCS, a notification made under Article 23 of NIS 2 satisfies the sectoral reporting obligation provided it contains the information the network code requires.[30] An impact assessment within 72 hours requires the operator to know, under pressure, which systems depend on the compromised one, which interfaces are live, which processes are business critical and who owns them. Where that dependency map does not exist in reliable form, the 24-hour and 72-hour statements are produced by reconstruction under time pressure. The regulatory record of European incidents is therefore, in part, a record of provisional inferences.

Fourth, audit verifies the declared scope. Under the NCCS, compliance by critical-impact entities is demonstrated through a national verification scheme or an independent third-party audit covering the verification scope at least once every three years.[31] An audit conducted against a perimeter defined by the auditee tests the implementation of controls within that perimeter. It does not test whether the perimeter itself is complete. Nor is it designed to do so. The same structural limit applies to the supply chain obligations under Article 21(3) of NIS 2 and Article 33 of the NCCS, since due diligence is performed on the suppliers the entity knows it has. The SBOM obligation under the CRA will improve component-level legibility from December 2027, but an SBOM issued by a manufacturer does not disclose the connectors, scripts and manual bridges through which an operator has integrated that product into its own estate.[32]

6. FROM THE UNKNOWN PERIMETER TO CRITICAL VULNERABILITY

The analysis so far has remained at the level of the operator and the regulator. Raised to the level of the State, the same problem becomes one of strategic knowledge rather than of scale.

a. The State inherits the operator's uncertainty

A State's assessment of its own critical-infrastructure vulnerability is assembled from information generated by several actors: infrastructure operators, sectoral regulators, national cybersecurity authorities, intelligence services, technology suppliers and sector-level monitoring organisations. The operator sits at the base of that building. It is the first layer through which national infrastructure is observed. Much of what the other actors know derives from information reported by operators.

European regulation makes this interdependence explicit. The NCCS establishes a common methodology and risk impact matrix at the Union level, together with the ECII and the thresholds used by competent authorities to identify high-impact and critical-impact entities. Member State assessments subsequently incorporate information relating to those entities and contribute to regional assessments. The resulting architecture is therefore not purely hierarchical. Union-level criteria shape what is assessed at entity and Member State level, while information generated at those levels adds to the broader assessment of cybersecurity risk.[33]

However, this circulation of information does not resolve uncertainty at its source. Where an operator's knowledge of its own installations is structurally incomplete, that uncertainty enters the assessment process with the information it provides. The national and regional picture may consequently represent not the infrastructure in its entirety but the portion of it that the reporting institutions have been able to observe. Aggregation does not necessarily correct that limitation and may conceal it behind the apparent completeness of the resulting assessment.

b. Declared, actual and discoverable infrastructure

The adversary is not bound by the declared perimeter. An organisation defends the systems it knows it has, whereas an adversary looks for the systems it can find. Three perimeters must therefore be distinguished rather than two. The declared infrastructure is what the organisation believes and states that it operates. The actual infrastructure is what is in fact running, including the informal integrations described in Part II. The discoverable infrastructure is what an external actor can identify, reach or exploit. These sets overlap without coinciding. The security consequence lies in the parts of the third perimeter that fall outside the first.

The two incidents examined in Part III illustrate the point from opposite directions. In the Ukrainian case, the operation did not begin by manipulating a generation or transmission asset directly. It developed through the surrounding corporate environment and moved progressively towards the control systems, which indicates that the adversary's effective perimeter was larger than the perimeter recognised by the organisations defending it.[34] In the Danish case, the campaign was identified through a shared sectoral sensor network rather than through the affected organisations' own documentation.[35] Knowledge about an infrastructure was in that instance produced outside the organisations that owned it. Both observations point the same way. The perimeter that matters operationally is not the one that has been declared.

The pattern is not confined to electricity. In its advice on the compromise of programmable logic controllers in water and other sectors, the United States Cybersecurity and Infrastructure Security Agency and its partners recorded devices reachable from the public internet and running default credentials. They advised owners to remove exposed controllers from the internet.[36] That advice can only be acted upon by an operator who knows that the device is exposed. In the documented cases, the adversary located the equipment through untargeted scanning before its owner was aware that it was locatable.

c. Second-order vulnerability and the confidence question

This suggests a distinction between two orders of vulnerability. The first order is the technical weakness itself: an unpatched controller, a flat network, or a credential that was never changed. The second order is the absence of knowledge that the weakness exists. The second is the harder of the two to manage because every conventional remedy presupposes identification. An asset that has not been identified cannot be prioritised, patched, monitored, segmented or written into contingency planning. For the same reason, it is also absent from the reports on which the national picture is built.

It does not follow that the State should aim at complete knowledge of its critical infrastructure. This article has argued that such knowledge is not available through declaration. An objective that cannot be met is of no use as a policy standard. What can be established is the confidence attaching to the representation. That requires both the State and the operators reporting to it to distinguish between (a) an asset known to be secure, (b) an asset whose security is uncertain, and (c) an asset whose existence or exposure is not known at all. The three categories carry different consequences and call for different responses. Collapsing them into a binary judgement of secure or insecure removes precisely the information a decision-maker needs.

Existing frameworks state the outcome without stating the confidence level. The NIST Cybersecurity Framework 2.0 formulates asset management as an outcome under which assets are identified and managed consistently with their relative importance to organisational objectives.

It situates that outcome within a governance function that makes cybersecurity risk a component of enterprise risk management.[37] It does not ask the organisation to state how completely it believes it has identified them. Requiring organisations to state how confident they are that their asset inventory is complete would make the remaining uncertainty visible to regulators and the State. Current reporting identifies what has been found, but does not reveal what may still be unknown.

Understood in this way, the unknown perimeter ceases to be a compliance difficulty for electricity companies and becomes a variable in national security. Uncertainty about the estate is not an accident of implementation that better tooling would remove, but a standing condition of complex infrastructure. The question for a State is whether that condition is measured and reported or left implicit in assessments that appear complete.

d. Why civilian infrastructure is the chosen target

The preceding sections have treated the adversary as a given. It is worth stating briefly why critical infrastructure attracts that attention because the answer bears on what a State is able to do in response.

Infrastructure serving a population is attractive precisely because interference with it produces political effects without producing the characteristics of armed conflict. Disruption of electricity or water imposes costs on civilians and tests national resilience. Operations of this kind are positioned deliberately below the threshold at which the law of armed conflict and the right of self-defence are engaged, a threshold that international lawyers continue to locate with difficulty in the cyber context.[38]

States have recognised the problem normatively without resolving it. The 2015 report of the United Nations Group of Governmental Experts records the recommendation that a State should not conduct or knowingly support activity that intentionally damages critical infrastructure or otherwise impairs its operation in the provision of services to the public, together with the recommendation that States take appropriate measures to protect their own. The General Assembly endorsed the report in the same year.[39] The norm is voluntary and non-binding. The record of the decade since suggests that it constrains conduct weakly.

The strategic literature explains the continuation of such operations. That explanation bears directly on the argument of this article. Coercion in cyberspace is generally assessed as a weak instrument taken on its own, since coercion depends on communicated threats, credibility and reassurance, none of which travels well through covert means.[40] What remains effective is not coercion but the accumulation of position, understood as the quiet establishment of access inside systems whose owners have not identified them as exposed. An unknown perimeter may therefore allow an adversary to establish a presence without detection or response, particularly where the affected asset has never been recorded by the State.

CONCLUSION

European cyber resilience law has correctly identified asset knowledge as foundational. NIS 2 makes it a minimum measure, the NCCS makes it a declared and audited perimeter, and the CRA makes it a product-level obligation. None of these instruments yet addresses the difficulty that large infrastructure operators face in producing that knowledge by declaration. Nor do they address the fact that the electricity industry is where the difficulty is most acute, given converged IT and OT environments, long-lived plant, constraints on active interrogation and accelerating decentralisation.

The regulatory question is whether the law will continue to accept declaration as evidence of knowledge or will begin to ask how that knowledge was produced and with what level of confidence. At State level, the same uncertainty becomes a strategic vulnerability where an adversary can identify and access infrastructure that remains outside the defender's known perimeter. A State's capacity to protect its critical infrastructure therefore depends on its ability to assess both the infrastructure itself and the limits of its knowledge of that infrastructure.

WORK CITED:

1. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union (NIS 2 Directive) [2022] OJ L333/80, art 21(2)(i).
2. Commission Delegated Regulation (EU) 2024/1366 of 11 March 2024 supplementing Regulation (EU) 2019/943 by establishing a network code on sector-specific rules for cybersecurity aspects of cross-border electricity flows, OJ L, 2024/1366, 24.5.2024, art 26(2).
3. Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements (Cyber Resilience Act), OJ L, 2024/2847, 20.11.2024, Annex I, Part II, point (1).
4. NIS 2 Directive (n 1) art 21(1) and (2). The ten minimum measures are technology-neutral and outcome-based.
5. ibid., art 20(1): management bodies approve the cybersecurity risk-management measures, oversee their implementation and *"can be held liable for infringements by the entities of that Article"*.
6. ibid., art 21(3), read with art 22(1) on Union-level coordinated security risk assessments of critical supply chains.
7. NCCS (n 2). Adopted 11 March 2024. Entry into force 13 June 2024. See EUR-Lex, ELI <http://data.europa.eu/eli/reg_del/2024/1366/oj> accessed 9 August 2026.
8. NCCS (n 2) art 24, on the identification of high-impact and critical-impact entities by national competent authorities on the basis of the Electricity Cybersecurity Impact Index. Provisional ECII thresholds were published by ENTSO-E and the EU DSO Entity: ENTSO-E, *Provisional Electricity Cybersecurity Impact Index* (2025).
9. NCCS (n 2) art 26(2). See also Emissions-EUETS, 'Network Code on Cybersecurity (NCCS)' (updated 27 January 2026) <https://emissions-euets.com/network-codes/network-code-on-cybersecurity> accessed 9 August 2026.
10. NCCS (n 2) arts 26(5) and 27.
11. ibid., arts 25 and 31.

12. ibid., art 33, in particular art 33(2) on minimum and advanced supply chain controls across the ICT lifecycle.
13. Cyber Resilience Act (n 3) Annex I, Part II, point (1), requiring a software bill of materials "in a commonly used and machine-readable format covering at the very least the top-level dependencies of the product".
14. European Commission, 'Cyber Resilience Act', Shaping Europe's Digital Future <https://digital-strategy.ec.europa.eu/en/policies/cyber-resilience-act> accessed 9 August 2026. Reporting obligations apply from 11 September 2026. The remainder of the Regulation applies from 11 December 2027.
15. Paul R Lawrence and Jay W Lorsch, Organization and Environment: Managing Differentiation and Integration (Harvard Business School Press 1967).
16. Michael E Porter, Competitive Advantage: Creating and Sustaining Superior Performance (Free Press 1985) ch 2.
17. World Economic Forum in collaboration with the University of Oxford, The Cyber Resilience Compass: Journeys Towards Resilience (White Paper, April 2025) 20.
18. International Energy Agency, Electricity 2026: Analysis and Forecast to 2030 (IEA 2026), Executive Summary.
19. Zoya Pourmirza, Mehmet Bozdal, Mohamad Khalil, Emily Judson and Sara Walker, 'Digital Transformation of Energy Systems: Technologies, Data, Governance and Cyber Security' (2026) IET Smart Grid (open access review, accepted 23 January 2026), section 7 and Conclusion.
20. Cybersecurity and Infrastructure Security Agency, 'Cyber-Attack Against Ukrainian Critical Infrastructure' (Alert IR-ALERT-H-16-056-01, 25 February 2016) <https://www.cisa.gov/news-events/ics-alerts/ir-alert-h-16-056-01> accessed 9 August 2026.
21. SektorCERT, The Attack Against Danish Critical Infrastructure (November 2023). See also Recorded Future News, 'Nearly two dozen Danish energy companies hacked through firewall bug in May' (13 November 2023) <https://therecord.media/danish-energy-companies-hacked-firewall-bug> accessed 9 August 2026.
22. Cyber Resilience Act (n 3) art 69(2): products with digital elements placed on the market before 11 December 2027 are subject to the Regulation "only if, from that date, those products are subject to a substantial modification".

23. Cyber Resilience Act (n 3) arts 69(3) and 14. The reporting obligations concerning actively exploited vulnerabilities and severe incidents extend to products placed on the market before 11 December 2027 and apply from 11 September 2026.
24. Pourmirza and others (n 19), noting that energy systems are *"safety-critical, time-sensitive and often constrained by legacy hardware, long equipment lifecycles and strict availability requirements"*.
25. International Energy Agency (n 18) 12.
26. Cyber Resilience Act (n 3) art 13(8), requiring the support period to reflect the length of time the product is expected to be in use and to be no shorter than five years unless the anticipated lifetime is shorter.
27. NCCS (n 2) art 26(2), which applies to all assets within the designated perimeters irrespective of the date on which they were placed on the market.
28. NIS 2 Directive (n 1) art 20(1).
29. ibid., arts 32(5) and 34(4).
30. ibid., art 23(4); NCCS (n 2) art 38(8).
31. NCCS (n 2) art 25(2)(e) and 31(1)–(4).
32. NIS 2 Directive (n 1) art 21(3); NCCS (n 2) art 33; Cyber Resilience Act (n 3) Annex I, Part II, point (1).
33. NCCS (n 2) arts 19 to 21, 26 and 27, establishing the tiered risk assessment architecture that runs from entity level through Member State level to the Union-wide assessment conducted by ENTSO-E and the EU DSO Entity.
34. CISA (n 20). See also NATO CCDCOE, *Cyber Law Toolkit,* 'Power grid cyberattack in Ukraine (2015)' <https://cyberlaw.ccdcoe.org/wiki/Power_grid_cyberattack_in_Ukraine_(2015)> accessed 20 August 2026. Published accounts differ on the number of substations affected and on whether the number of consumers is best given as approximately 225,000 or 230,000. The figures used above follow the CISA alert.
35. SektorCERT (n 21).
36. Cybersecurity and Infrastructure Security Agency and partners, 'IRGC-Affiliated Cyber Actors Exploit PLCs in Multiple Sectors, Including US Water and Wastewater Systems Facilities' (Joint Cybersecurity Advisory AA23-335A, 1 December 2023) <https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-335a> accessed 20 August 2026.

37. National Institute of Standards and Technology, The NIST Cybersecurity Framework (CSF) 2.0 (NIST CSWP 29, 26 February 2024) Appendix A, ID.AM and the Govern function.
38. Michael N Schmitt (ed), Tallinn Manual 2.0 on the International Law Applicable to Cyber Operations (2nd edn, Cambridge University Press 2017), on the difficulty of locating the threshold of the use of force and of armed attack in the cyber context.
39. United Nations General Assembly, Report of the Group of Governmental Experts on Developments in the Field of Information and Telecommunications in the Context of International Security (UN Doc A/70/174, 22 July 2015) para 13(f) and (g). The report was endorsed by UNGA Res 70/237 (December 2015).
40. Erica D Borghard and Shawn W Lonergan, 'The Logic of Coercion in Cyberspace' (2017) 26 Security Studies 452.