OEM EU Readiness

Eighteen pieces on cybersecurity, financing, and architectural demarcation for non-EU OEMs engaging European-financed renewable energy projects.

Version 1.0.0 — 2026-05-18
DOI: 10.5281/zenodo.20268560

Eighteen pieces on cybersecurity, financing, and architectural demarcation for non-EU OEMs engaging European-financed renewable energy projects.

Rajesh Khanikar
ORCID: 0009-0008-8976-4491
Version 1.0.0 — 2026-05-18

DOI: 10.5281/zenodo.20268560 (concept)
Canonical: https://khanikar.com/series/oem-eu-readiness/

To cite this pack:

> Khanikar, R. (2026). *OEM EU Readiness* (Version 1.0.0). https://doi.org/10.5281/zenodo.20268560. Licensed under CC BY 4.0.

This pack is offered as editorial guidance for technical and procurement audiences. It does not constitute legal, financial, regulatory, or professional advice; the author is not a lawyer, an auditor, or a financial adviser. The content is provided without warranty of any kind, express or implied — no guarantee of accuracy, completeness, currency, or fitness for any particular procurement, project, or jurisdiction. Readers must verify against the primary sources and consult qualified professionals before acting on anything in this pack. EU regulations referenced are linked to their EUR-Lex ELI URLs, which are the canonical permalinks; texts may have been amended since the publication date above. The author accepts no responsibility or liability for any decision, action, or omission made in reliance on this content.

Corrections, citation-rot reports, and procurement-conversation experiences from new venues are welcome at linkedin.com/in/rajeshkhanikar.

Contents

When the money comes from Europe, the rulebook travels with it

15 May 2026 · 4 min read · #compliance #security #industrial #oem-eu-readiness

A renewable energy plant in North Africa or the Middle East. Around a hundred and fifty megawatts. The project sponsor is an EU-headquartered independent power producer. The lenders are a syndicate of European banks. The successful equipment bidder is a manufacturer with a strong track record across Asia and the Belt and Road region, bidding aggressively on capital cost, offering an integrated service package, confident in their proposal.

The first technical meeting goes well until the conversation turns to cybersecurity.

The bidder's engineering team is well prepared on the things they have always been asked about — redundancy, availability, mean time between failures, the architecture of their condition-monitoring system. They are less prepared when the project sponsor's architect asks about their vulnerability disclosure programme, the public list of security advisories for their products, the certificate from an independent auditor attesting their development process against an international standard most of the room knows only by its number.

The questions are unfamiliar. The expected answers are missing. The bidder leaves the meeting with a list of things to investigate, slightly off balance, wondering why a project physically located in Africa, financed for a customer based in Europe, operated in a non-European regulatory space, has just been measured against rules written in Brussels.

Here is the short answer to that question.

The money is European. The borrower is European. The borrower's auditors, insurers, regulators and shareholders are European. Each of those parties has obligations of its own, written into the rules of the jurisdictions where they live. Those obligations travel down the chain — from the regulator to the bank, from the bank to the borrower, from the borrower to the supplier — as contractual requirements, even when the project itself sits in a country those regulators have no direct reach into.

The rules of the European Union, on cybersecurity, on data protection, on supply chain transparency, on the secure development of products with digital elements, have over the last several years become a coherent stack. They apply, directly, to manufacturers placing products on the European market. They apply, indirectly but firmly, to anyone who supplies into a project whose ownership, financing, or operation passes through European hands.

This series is for those suppliers.

What follows is a seventeen-piece walk-through of what a manufacturer bidding into an EU-linked renewable energy project should expect to find on the table. Not what we hope they will do. Not what would be ideal. What is already, today, a normal condition of doing business when the financing chain reaches back into Europe.

Some of it will be unfamiliar. The architecture for remote access into the plant. The treatment of the communication network as belonging to the operator, not the supplier. The boundary at the substation. The cellular modem on the controller board. The public-facing security advisories page. The bill of materials of the firmware. The cryptography baseline. The identity and access model. The patch delivery contract. The cross-border data flows. The sanctions disclosure. The mismatch between a five-year service contract and a twenty-five-year support commitment.

Each of these will get its own article. Each will explain what the expectation is, why the expectation exists, and what a manufacturer can do to meet it. The first piece after this one is a fifteen-minute walk through the regulatory stack itself — the names of the laws, what each requires, and how each lands as a clause in a lender's term sheet. The final piece is a procurement timeline matrix showing where each conversation belongs, from the request for information through to mid-life review.

The series is written in the conviction that most non-European suppliers are losing competitive ground in EU-linked projects not because their products are inferior, but because nobody has sat them down and told them what is being measured against. The list is long but not difficult. The work is real but not prohibitive. The advantage on capital cost that a non-European manufacturer brings into a North African or Middle Eastern bid is, in most cases, more than enough to absorb the cost of meeting these requirements — but only if the requirements are understood before the bid lands rather than after.

The cheapest way to do this work is at proposal stage. The most expensive way is after preferred bidder selection, when the gaps surface during lender due diligence and the project programme is already committed. Most of what follows is, in effect, a guide to doing the work at the cheap end of that curve.

The next article in the series sets out the legal substance. The sixteen articles after that translate it into the engineering, commercial and operational decisions a manufacturer will need to make.

This is the opening anchor of an eighteen-part series on EU readiness for non-European suppliers. If a particular procurement conversation

has caught your team off-guard and you would like it covered in the series, LinkedIn is the way to suggest it.

The regulatory stack, in fifteen minutes

15 May 2026 · 16 min read · #compliance #security #industrial #oem-eu-readiness

This is the reference piece. The sixteen articles that follow it will repeatedly mention five pieces of European law and one framework that is technically not law at all. Each is covered here briefly, with links to the primary sources, and with a single thread running through every section: how the regulation in question lands, in practice, as a clause in the lender's term sheet.

The reason for that thread is operational. A manufacturer reading the regulation directly will find it abstract, lengthy, and addressed to parties other than themselves. A manufacturer reading the lender's draft loan agreement will find specific conditions, deadlines and deliverables that look unfamiliar but are not, in fact, the lender's invention. They are the regulation, translated downstream by lawyers whose job it is to make sure their bank does not breach its own obligations.

Five regulatory instruments, one industry framework. The Cyber Resilience Act, the NIS2 Directive, the cross-border data provisions of the General Data Protection Regulation, the EU Taxonomy and the wider sustainable finance regime, and the Equator Principles. The order below is roughly the order in which they appear in a typical project's compliance work — the product-level requirements first, the operator's obligations next, the data transfer questions when the architecture is on the table, and the lender's framing visible throughout.

The Cyber Resilience Act

Formally Regulation (EU) 2024/2847. Entered into force on 10 December 2024. Vulnerability reporting obligations apply from 11 September 2026. Full application from 11 December 2027.

The Cyber Resilience Act is the regulation that does most of the work in this series, because it is the only one of the five that addresses the manufacturer directly. The other instruments speak to operators, lenders, or data controllers, and reach the manufacturer through contractual flow-down. The Cyber Resilience Act speaks to the manufacturer. The CRA applicability post covers the scope question in more detail.

It applies to "products with digital elements" placed on the EU market. A wind turbine controller is a product with digital elements. A solar inverter is. A battery management system, an OT switch with embedded firmware, an OPC UA gateway, an HMI, a SCADA workstation, a Modbus-to-IEC 60870 protocol converter — all products with digital elements. The regulation applies regardless of where the product is manufactured, regardless of who buys it, as long as the product is placed on the EU market.

The key obligations are five. First, secure-by-design and secure-by-default development: manufacturers must conduct a cybersecurity risk assessment before placing the product on the market, address known vulnerabilities, and document the technical decisions taken. Annex I of the regulation lists the essential requirements; the harmonised standards that will operationalise those requirements are expected to reference IEC 62443 for industrial products.

Second, a declared support period. The manufacturer must publicly state how long the product will receive security updates. Article 13(8) sets the floor: the support period shall not be shorter than five years, except where the expected period of use of the product is shorter. For

industrial controllers and embedded systems in plants designed to operate for twenty to thirty years, the five-year floor is far short of what the regulation's intent requires — the declared support period must reflect actual expected use, not the statutory minimum.

Third, vulnerability disclosure and patch obligations. The manufacturer must operate a coordinated vulnerability disclosure programme, publish information about known vulnerabilities and their patches, and provide security updates free of charge within the declared support period. Article 14 requires notification of actively exploited vulnerabilities simultaneously to ENISA and the designated CSIRT coordinator without undue delay and in any event within 24 hours of becoming aware of the vulnerability, follow-up reporting within 72 hours, and final reports within 14 days of a corrective measure being available.

Fourth, conformity assessment and CE marking. Most products are subject to self-assessment by the manufacturer; "important" and "critical" categories require third-party conformity assessment by a notified body. The CE mark on the product signifies conformity with the essential requirements.

Fifth, technical documentation, including a software bill of materials made available to market surveillance authorities. The bill of materials is the inventory that underpins the vulnerability handling obligation; a manufacturer that does not know what is in their firmware cannot meaningfully claim to handle vulnerabilities in it.

Geographic reach is straightforward. If the product is placed on the EU market — even once — the manufacturer is in scope for that product worldwide. A turbine model sold into a Spanish offshore project is in scope, and the same model deployed into a North African project inherits the disclosure infrastructure, the bill of materials, the CE marking and the support period commitment for free. A model never

placed on the EU market technically sits outside the regulation, but the lender's term sheet will require equivalence regardless.

In a lender's term sheet, the Cyber Resilience Act appears as conditions precedent and conditions subsequent. Before drawdown: evidence of CE marking for in-scope products, the manufacturer's declared support period stated in writing, the URL of the vulnerability disclosure page, the bill of materials delivery commitment. During operation: notification of any actively exploited vulnerability under the same 24/72-hour cadence, the right to audit the manufacturer's vulnerability handling programme, the right to terminate or step in if the manufacturer fails to provide patches within the support window.

Primary source: Cyber Resilience Act — European Commission . Regulation text: Regulation (EU) 2024/2847 .

The NIS2 Directive

Formally Directive (EU) 2022/2555, the second Network and Information Security Directive. Replaced the original NIS Directive in January 2023. Member states were required to transpose it into national law by 17 October 2024. Transposition has been uneven across the bloc, but the substantive obligations are now in force across most of the EU.

NIS2 addresses operators, not manufacturers. For the project sponsor — the renewable energy company operating the plant — NIS2 is the most consequential of the five instruments, because it imposes direct obligations on the sponsor as an "essential entity" in the energy sector. Electricity producers above a defined threshold are essential entities by default, and most EU-headquartered renewable energy companies operating utility-scale assets are in scope. The NIS2 applicability post covers the scope question in more detail.

The obligations have several components. Risk management measures, listed at Article 21, include incident handling, business conti-

nuity, supply chain security, security in network and information systems acquisition, vulnerability handling and disclosure, basic cyber hygiene practices and training, cryptography, human resources security, access control, asset management, and multi-factor authentication. Governance, at Article 20, places direct accountability on the management body — directors must approve risk-management measures and oversee their implementation, and can be banned from executive functions under Article 32(6) if the entity persists in non-compliance. Incident reporting, at Article 23, requires early warning to the national CSIRT within 24 hours of a significant incident, an incident notification within 72 hours, and a final report within one month. The NIS2-to-IEC-62443 mapping post walks through each Article 21 measure against the corresponding 62443 clauses.

The manufacturer rarely meets NIS2 directly. The manufacturer meets the operator's flow-down of NIS2 obligations, particularly the supply chain security clause. Article 21(2)(d) requires essential entities to address security in their supply chains, including assessing the cybersecurity practices of direct suppliers. The operator is contractually required to push these assessments down to the supplier, document the results, remediate gaps, and report on supply chain security to the competent authority. A turbine manufacturer becomes part of that supply chain, and the assessments flow through procurement.

Geographic reach is again partly contractual. NIS2 covers operators of services within the EU. An EU-headquartered IPP operating a plant in North Africa is operating that asset outside the EU, but the IPP's group governance, audit, board oversight, and consolidated reporting obligations under NIS2 reach the asset regardless of geography. A breach in the North African plant becomes a board-level event in the IPP's home capital, with the same reporting cadence and the same personal liability exposure for directors.

In a lender's term sheet, NIS2 appears as governance representations and warranties, ongoing covenants, and information rights. Representations and warranties: that the borrower has an information security management system, that it has identified its essential supply chain dependencies, that it has conducted a cybersecurity risk assessment for the project. Covenants: that the borrower will maintain those systems, will notify the lender of significant incidents under a defined cadence (usually mirroring NIS2 or stricter), will permit the lender's technical adviser to audit. Information rights: copies of incident reports, audit findings, remediation plans.

Primary source: Directive (EU) 2022/2555 . Implementation overview: NIS2 Directive — European Commission .

The GDPR, Articles 44–49

The General Data Protection Regulation — Regulation (EU) 2016/679 — has been in force since 25 May 2018. The articles relevant to this series are not the well-known ones about consent or data subject rights. They are Articles 44 to 49, which govern transfers of personal data to countries outside the European Economic Area.

The chain of logic is short. The GDPR applies to any controller or processor established in the EU. An EU-headquartered IPP operating a plant in North Africa is established in the EU, and any personal data it processes — workforce credentials, badge logs, CCTV around the control room, identifiable telemetry — is within scope wherever the data physically sits. When that data flows from the plant to a manufacturer's cloud for condition monitoring, performance analytics, or remote support, the flow is a "transfer to a third country" under Article 44.

Article 45 permits transfers to countries the European Commission has adjudged to provide an adequate level of data protection. The list includes the United Kingdom, Switzerland, Japan, the Republic of Ko-

rea, New Zealand, Canada (commercial organisations), Israel, Argentina, Uruguay, the Faroe Islands, Guernsey, the Isle of Man, Jersey, Andorra, and — under the EU–US Data Privacy Framework adopted 10 July 2023 — the United States for recipients on the DPF list with valid certification. The major manufacturing jurisdictions for industrial control equipment outside the DPF perimeter are not on it.

Article 46 permits transfers without an adequacy decision if appropriate safeguards are in place. The most common safeguard is the Standard Contractual Clauses, updated by the Commission in June 2021. But the Standard Contractual Clauses alone are not enough after the Schrems II judgment.

Schrems II — Court of Justice of the European Union case C-311/18, decided 16 July 2020 — invalidated the EU-US Privacy Shield and held that controllers using the Standard Contractual Clauses must conduct a transfer impact assessment for the destination country. If the law of the destination country allows public authorities to access transferred data in ways that exceed what is necessary and proportionate under EU standards, the Standard Contractual Clauses do not on their own provide adequate protection and supplementary measures are required. The supplementary measures must be technical (encryption with keys held only in the EU, pseudonymisation, split processing) or organisational, and must close the gap identified in the transfer impact assessment.

For some destination countries — including the home jurisdiction of several major industrial equipment manufacturers — the gap identified by Schrems II-style analysis is not practically closable. National security and intelligence law regimes in those countries provide access to data held in their territory in ways that exceed what is necessary and proportionate under EU standards, and no technical measure short of refusing the transfer entirely satisfies the test.

For the project, this means architectural choices have legal consequences. Operational data and condition-monitoring telemetry that lands in a manufacturer's home-country cloud may be unlawful under the GDPR even with Standard Contractual Clauses in place. Personal data in particular — engineers' identities, access logs, video — must either stay in the European Economic Area, transit through an adequacy-decision country, or be processed in such a way that the manufacturer cannot in practice receive personal data at all.

In a lender's term sheet, Articles 44–49 appear as data flow representations and architectural commitments. Representations: that the borrower has identified all cross-border data flows, conducted transfer impact assessments where required, implemented appropriate safeguards. Commitments: that the architecture as built will keep personal data within agreed jurisdictions, that telemetry to the manufacturer will not include identifiable personal data without explicit derogation, that any change to the data flow architecture requires lender consent.

Primary sources: Regulation (EU) 2016/679 . Schrems II judgment: CJEU C-311/18 .

The sustainable finance regime

The European Union's sustainable finance framework is the indirect path by which cybersecurity reaches projects that no other instrument directly covers. None of its three main pillars — the Taxonomy Regulation, the Corporate Sustainability Reporting Directive, or the Sustainable Finance Disclosure Regulation — names cybersecurity in the way the Cyber Resilience Act or NIS2 do. But all three operate as the framework within which the lender's environmental, social and governance due diligence is conducted, and cybersecurity has migrated firmly into the governance category over the last five years.

The Taxonomy Regulation, Regulation (EU) 2020/852, defines when an economic activity is environmentally sustainable. For renewable energy projects, the substantial contribution criteria for climate change mitigation are relatively easy to meet — a wind farm or solar plant contributes substantially almost by definition. The harder tests are the "do no significant harm" criteria across the other five environmental objectives, and the minimum safeguards under Article 18, which require alignment with the OECD Guidelines for Multinational Enterprises and the United Nations Guiding Principles on Business and Human Rights. The minimum safeguards are the gateway through which broader governance expectations — including the management systems expectations that increasingly include cybersecurity — enter the Taxonomy assessment.

The Corporate Sustainability Reporting Directive, Directive (EU) 2022/2464, requires in-scope companies to report against the European Sustainability Reporting Standards. ESRS G1 (business conduct) and aspects of ESRS S1 (own workforce) and S4 (consumers and end users) bring information security and data protection within mandatory disclosure. A lender financing a project for a CSRD-scoped sponsor is financing an asset whose cyber posture will appear in the sponsor's consolidated sustainability statement. That visibility creates downstream procurement discipline.

The Sustainable Finance Disclosure Regulation, Regulation (EU) 2019/2088, applies to financial market participants, including lenders themselves, requiring them to disclose how their products consider sustainability risks and adverse impacts. Cyber incident exposure is increasingly identified as a sustainability risk in lender frameworks, and the principal adverse impact indicators that lenders report against include governance failures that often have a cyber component.

The practical effect of the sustainable finance regime on a non-EU manufacturer is not a specific obligation. It is a tone. The lender's term sheet, its environmental and social action plan, its ongoing reporting requirements, all sit within a framework that expects the project to be governed to European standards on management systems, supply chain due diligence, and operational resilience. Cybersecurity sits inside that envelope. When the lender's adviser asks for evidence of the manufacturer's information security management system, or for the supplier's policy on responsible business conduct, the question is grounded in this regime even if the lender does not cite it.

In a lender's term sheet, the sustainable finance regime appears as the framing of the entire environmental and social action plan, the basis for the borrower's reporting covenants, and the justification for the lender's right to engage technical and ESG advisers throughout the life of the loan.

Primary sources: Taxonomy Regulation (EU) 2020/852 ; CSRD — Directive (EU) 2022/2464 ; SFDR — Regulation (EU) 2019/2088 .

The Equator Principles

The Equator Principles are not law. They are a voluntary risk management framework adopted by financial institutions for determining, assessing and managing environmental and social risk in project finance. The fourth iteration, EP4, came into effect on 1 October 2020 and remains the current version. As of early 2026, approximately 128 financial institutions across 38 countries are signatories, covering the majority of international project finance debt in emerging and developed markets. The Equator Principles Association was succeeded by Equator Principles Limited (legal entity from 1 January 2024); the steering-committee governance is unchanged.

For projects whose financing structure qualifies as project finance — most renewable energy independent power producers do — the

lender's signature on the Equator Principles is the operational mechanism that imports the framework into the project. A signatory institution will not provide finance to a project that does not comply with the Principles. The Principles, in turn, reference the International Finance Corporation's Performance Standards on Environmental and Social Sustainability as the substantive baseline.

The Performance Standards are eight in number. Performance Standard 1 (assessment and management of environmental and social risks and impacts) is the foundational one for cybersecurity, because it requires the client to establish and maintain an environmental and social management system commensurate with the project's risks. Cyber risk has become an explicit category within such management systems over the last several years, particularly for energy and infrastructure projects.

The framework also embeds a stakeholder engagement requirement (Principle 5), a grievance mechanism requirement (Principle 6), and an independent review requirement (Principle 7). For Category A and high-risk Category B projects — which most utility-scale renewable energy projects are — an independent environmental and social consultant is appointed to review the borrower's compliance with the Principles and the underlying Performance Standards. The consultant's review increasingly includes a cyber risk assessment, particularly where the project depends on remote operations, manufacturer service connectivity, and cross-border data flows.

For a non-EU manufacturer, the Equator Principles appear in two places. First, in the project's environmental and social action plan, which lists the specific cybersecurity-related commitments the borrower has made to the lenders. Many of those commitments cascade to suppliers as technical specifications and contract conditions. Second, in the independent reviewer's report, which may flag the manu-

facturer's cyber posture as an outstanding action item before disbursement, or as a condition subsequent during operation.

In a lender's term sheet, the Equator Principles appear as the framing for the environmental and social action plan, the basis for the appointment of the independent reviewer, and the reporting and audit rights that survive the construction phase into operations.

Primary source: The Equator Principles .

What this stack actually says

Five instruments, one framework. The Cyber Resilience Act for the product. NIS2 for the operator. The GDPR's transfer provisions for the data flow. The sustainable finance regime for the governance envelope. The Equator Principles for the financing envelope. None of them, read in isolation, captures the full picture. Read together, they describe a single, coherent expectation: that the project is governed and operated to European standards, that the products in it are secure by design and supported through their lifetime, that personal data does not leak across borders the European Union does not regard as safe, and that the lender has visibility and rights throughout.

A non-EU manufacturer cannot make all of this go away. The regulations are not negotiable. The lender's term sheet, in any meaningful EU-financed project today, will reflect them.

But the regulations are also not as forbidding as they sometimes appear in the first conversation. Most of the work is procedural and architectural. The product changes that are required — the bill of materials, the disclosure page, the cryptographic baseline, the support period commitment — are achievable within the normal product development cycle once they are understood as priorities. The architectural changes — the remote access model, the data flow design, the network demarcation — are choices a manufacturer is well placed to

influence at proposal stage, when the bid still allows trade-offs to be made.

The next article in the series begins the substantive walk-through with the most architecturally consequential of those choices: the treatment of the communication network as belonging to the operator, not the supplier.

This article reflects the regulatory state at publication. Subsequent CRA implementing acts, NIS2 transposition activity, evolution of the EU sustainable finance regime, or revision of the Equator Principles may shift specific obligations described above. Specific transactions should be reviewed by qualified legal counsel rather than against this article. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

The communication network is not yours to design

15 May 2026 · 9 min read · #compliance #security #industrial #oem-eu-readiness

The manufacturer's network engineer arrives at the early-design meeting with a topology drawing. It shows the turbine controllers, the SCADA workstation, the historian, the engineering workstation, and the firewall to the outside world. The IP addresses are pre-assigned. The VLANs are labelled. The redundant uplink connects to a switch marked "owner network", and at that switch the drawing politely stops.

The drawing is, by its own conventions, complete. It shows everything from the sensor in the blade to the boundary of the operator's network. The conversation that follows assumes the drawing is the basis for the network design — that the manufacturer specifies, the operator implements what the manufacturer specifies, and the boundary is where the two diagrams meet.

That assumption is the source of more friction in early project meetings than any other single architectural question.

The plant communication network is not the manufacturer's to design. The local area networks inside the plant, the segmentation between zones, the firewalls between layers, the industrial demilitarised zone that mediates upstream connectivity, the routing to the substation and to corporate, the IP address space, the VLAN allocation, the default gateways, the spanning-tree topology, the time synchronisation hierarchy — all of it is the operator's responsibility. The manufacturer specifies what their equipment needs. The operator specifies how that need is met.

This is not a courtesy. It is a structural consequence of the regulatory and security framework the project sits inside.

Why this separation exists

Three reasons make the role separation non-negotiable.

The first is governance. Under the NIS2 Directive , the operator is the essential entity. The operator is the party with statutory obligations on risk management, incident handling, supply chain security and reporting. The operator's regulator does not accept "the manufacturer designed and operates the network" as an answer when something fails. A network that the operator cannot describe, cannot monitor, cannot audit and cannot change without manufacturer approval is a governance failure on the operator's side, regardless of how technically well-designed it might be.

The second is the security architecture itself. IEC 62443 organises an industrial control system into zones — groupings of assets that share a common security level — separated by conduits, the controlled communication paths between them. The asset owner is responsible for the zone and conduit design (62443-3-2) and the system-level security requirements (62443-3-3) . The manufacturer, as supplier and L0/L1 integrator, is responsible for the components within their zone (62443-4-1 for the development process, 62443-4-2 for the components themselves) . The architecture is explicit about where each party's design authority begins and ends, and the plant communication network — particularly the network above the manufacturer's process zone — sits firmly on the operator's side of that line.

The third is the multi-vendor reality of modern renewable energy plants. A typical utility-scale project includes a turbine or solar inverter manufacturer, a battery energy storage supplier with its own controls, a SCADA platform vendor whose software may come from a third party, a substation automation vendor, a protection relay vendor,

sometimes a separate condition-monitoring vendor, and a meteorological mast supplier. Each cannot bring their own network. Each connects to a shared infrastructure designed and operated by the asset owner. A manufacturer who arrives expecting to design the network they will sit on has not yet accommodated the fact that they are one tenant among several.

What good vendor specification looks like

What a manufacturer should specify is precise and brief. The equipment, by model and quantity. The communication protocols, by standard. The transport requirements — TCP or UDP, port numbers, expected message rates, latency tolerance, packet loss tolerance, jitter sensitivity. The physical layer — copper or fibre, single-mode or multimode, port count per device, connector type, distance constraints. The redundancy posture — single-homed, dual-homed with application-level failover, parallel redundancy protocol, high-availability seamless redundancy, rapid spanning tree with stated convergence targets. The time synchronisation requirement, in terms of accuracy and protocol (NTP, PTP, IRIG-B).

That is the entirety of what the manufacturer needs to say. The operator's network team takes that specification and delivers a network that meets it.

What a manufacturer should not specify is the implementation. They should not state that their controller requires the 192.168.10.0/24 subnet. They should not state that the redundant LAN must use VLAN 200. They should not state that the gateway IP must be 192.168.10.1. They should not state that the time source must be a specific NTP server address embedded in firmware. They should not assume that their device will be on the same broadcast domain as any other device they communicate with. They should not assume the firewall rule between their device and the historian will be open in

both directions for a port range; it will not. It will be open from the device to the historian, for the specific port, with the specific source and destination addresses, and any deviation from the specified behaviour will be dropped.

A device that comes off the production line with hard-coded subnets, embedded gateway assumptions or fixed IP addresses that cannot be configured at deployment is a device with a delivery defect. It will be returned. The manufacturer will be asked to provide firmware that accepts IP configuration parameters from the operator's deployment process. If that firmware does not exist, the manufacturer's bid does not survive the technical evaluation.

Redundancy and the assumed topology

The redundancy question deserves its own treatment because it is where the assumed topology most quietly collides with the operator's design.

A manufacturer's redundancy requirement, stated correctly, looks like this: "Controller A requires two independent network interfaces with failover between them, recovery time below 200 milliseconds, no loss of in-flight transactions during a single link failure." That is a clean specification. The operator can deliver it through Parallel Redundancy Protocol or High-availability Seamless Redundancy under IEC 62439-3, through Rapid Spanning Tree with appropriate timer tuning, through link aggregation with LACP, through Multiple Spanning Tree with engineered convergence, or through application-managed failover where the controller maintains parallel connections on both interfaces. There are many implementations; the manufacturer cares that the requirement is met, not how.

A manufacturer's redundancy requirement, stated incorrectly, looks like this: "Controller A requires LAN 1 on subnet 192.168.10.0/24 and LAN 2 on subnet 192.168.20.0/24, both with the same default gate-

way, failover triggered by loss of ICMP response from the gateway, return to primary after 30 seconds of stable response."

That is not a requirement. It is an implementation, and it is an implementation that assumes the operator's network can accommodate the manufacturer's mental model of redundancy. Most plant networks cannot, because they run protocol-level redundancy at the switch fabric rather than at the device, or they run application-level redundancy where the device opens connections on both interfaces and uses whichever responds, or they use addressing schemes — overlapping management VLANs, non-routable maintenance networks, unique-local IPv6 — that the manufacturer's static-IP assumption simply breaks.

The architectural principle is straightforward. The manufacturer assumes that their device will receive an IP address. From which CIDR block, with which subnet mask, with which default gateway, on which VLAN — none of these are the manufacturer's to specify. Each interface configurable at deployment, each tolerant of the addressing scheme the operator chooses, each functioning correctly regardless of whether the two interfaces sit on the same subnet, on different subnets, on different VLANs or on different physical networks entirely. The device's redundancy logic must work in all of those conditions, because in different plants in different jurisdictions designed by different network teams, all of those conditions will occur.

The visibility question

The visibility question is the hardest part of this conversation for manufacturers to accept, because it feels like an information vacuum. The manufacturer is told that their equipment will be deployed in a network whose topology, addressing scheme, segmentation and firewall rule base they will not see in any meaningful detail.

This is not gatekeeping. It is the security framework operating as intended.

The plant network topology is a sensitive asset. It maps the attack surface. It identifies which assets are reachable from which other assets, which traffic flows the firewalls permit, which segments are isolated for safety reasons, which paths a sophisticated attacker would have to traverse to reach the critical zones. Each additional party that holds the topology is a party whose own information security posture becomes a path to the plant. The principle of least privilege, applied to design information rather than to access credentials, says that the manufacturer should see only what the manufacturer needs to see.

What the manufacturer needs to see is what they actually receive. The IP addresses their devices will use. The gateways their devices will route through. The destination addresses and port numbers their devices will reach. The protocols they will speak on each interface. The credentials for their service accounts. The diagnostic interface they may query for the health of their own equipment. They are not given the diagram of how those endpoints are reached internally. They are not given the list of other devices on the same VLAN. They are not given the firewall rule base above the conduit between their zone and the next. That information is not theirs.

What the manufacturer is given in addition, at the conduit boundary, is a specification of what their equipment must support: which protocols are allowed across the conduit, in which directions, with which authentication and which logging. The conduit specification is the contract. Everything beyond the conduit is the operator's domain.

What this means at proposal stage

At proposal stage, this discipline pays off rapidly. A bid that specifies the equipment, the protocols, the port numbers, the bandwidth enve-

lope, the latency budget, the redundancy posture and the time synchronisation requirement — and stops there — is a bid the operator's network architects can work with. A bid that specifies the equipment plus a full assumed topology with hard-coded addresses, embedded VLAN identifiers and required subnet masks is a bid that will be returned with comments asking the manufacturer to remove the topology and restate the requirements.

The principle is not new. It is how interfaces have always worked between independent engineering disciplines. The manufacturer who treats the network specification as an interface contract — what I need from you, what you need from me, no more — finds the architectural conversation much shorter than the manufacturer who arrives with a topology drawing.

The deeper habit to break is the one imported from markets where the manufacturer was also the integrator, the operator, and the network designer. In an EU-financed project, those roles are separated by design. The plant network is the operator's, and the operator's alone. The manufacturer's job is to specify what they need from it and to deliver equipment that works inside whatever the operator builds.

The next article moves one step outward, to the boundary where the plant ends and the substation begins — and to the engineering discipline on the other side of that fence, which is not the operator's either.

This article reflects the regulatory and standards landscape at publication. References to IEC 62443 and IEC 62439 may be superseded by revisions of those standards; NIS2 transposition continues to evolve across member states. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

The substation is a foreign country

15 May 2026 · 8 min read · #compliance #security #industrial #oem-eu-readiness

The manufacturer's lead control engineer asks, perfectly reasonably, for the single-line diagram of the 220 kV substation. They want to understand the busbar arrangement, the transformer impedances, the protection scheme, the timing of the reclosers, the disturbance recorder configuration. They are designing a wind farm controller and want to model how the substation will respond to faults so the turbine controller can ride through them appropriately.

The asset owner's substation team declines.

Not impolitely. Not as an act of obstruction. They decline because the document the manufacturer has asked for sits inside a different engineering discipline, owned by different people, subject to different standards, regulated by a different authority, and shared on different terms — and because the manufacturer does not need it to do their job.

The substation is a foreign country. The wind or solar plant ends at a defined electrical and informational boundary, and beyond that boundary the engineering belongs to another discipline. The grid operator sets the rules through the connection agreement. The protection and control engineers set the configuration through the protection coordination study. The substation cybersecurity team — sometimes the asset owner's, sometimes the transmission system operator's, sometimes a third party — sets the security architecture inside the fence. The manufacturer specifies what comes across the boundary, in which protocol, with which data points, at what update rate. Everything beyond the boundary is opaque by design.

This applies whether the substation in question is the on-site collection substation built specifically for the project, or the transmission substation connecting the plant to the grid. The physical location does not determine the discipline. A 33 kV collection substation sitting fifty metres from the nearest wind turbine is still substation engineering, still governed by protection coordination, still subject to the grid code, still inside a different design authority than the turbines that feed it.

Why the substation is its own jurisdiction

Several reasons converge to make this separation non-negotiable.

Grid code compliance is national. The grid code is set by the national transmission system operator and the energy regulator. It defines how plants must behave at the point of connection — fault ride-through capability, frequency response, reactive power range, voltage regulation, harmonic emissions, communication and telemetry to the control centre. Compliance is demonstrated at the substation interface and witnessed by the grid operator. The manufacturer contributes to grid code compliance through the behaviour of the equipment they deliver, but the demonstration is conducted at, and on the substation side of, the boundary.

Protection coordination is a separate engineering discipline. Protection settings are calculated by P&C engineers who model the entire network — not only the plant, but upstream lines, neighbouring substations, and the grid behind them. Settings interact with reclosers, disturbance recorders, breaker failure protection and busbar protection schemes that have been engineered to coordinate across the whole grid section. A manufacturer who wants to alter, or even to fully understand, the protection scheme is reaching into engineering that is neither under the asset owner's contractual control nor under the grid operator's design authority.

Cybersecurity inside the substation is its own zoning problem. Substations operate under IEC 61850 station-bus and process-bus architectures, frequently with intelligent electronic devices from different vendors than the plant SCADA, their own time synchronisation infrastructure, and their own security framework under IEC 62351. The cybersecurity zoning inside a substation is a separate design exercise from the plant's IEC 62443 zoning, conducted by different engineers, with its own threat model and its own conformance evidence.

Physical and informational security is separately regulated. In many jurisdictions substations are protected under critical infrastructure or counter-terrorism legislation that imposes its own personnel vetting, physical access control and reporting obligations. The information that describes a substation — the single-line, the protection scheme, the physical layout, the list of personnel with access credentials — is regulated information. Who holds it, how it is transmitted, and on what authority it is shared, are matters governed by law rather than by contract.

What the manufacturer specifies at the boundary

At the boundary the manufacturer specifies a fairly precise set of things.

The protocol used to exchange data. For most modern projects this is IEC 60870-5-104 for telecontrol to and from the grid control centre, IEC 61850 MMS or GOOSE for higher-bandwidth integration with substation automation, occasionally DNP3 in markets where it remains dominant, and IEC 61400-25 for wind-specific monitoring extensions over the 61850 framework.

The data points exchanged — analogue measurements, status indications, control commands, sequence-of-events records, fault records — usually documented in an interface control document that lists every

point, its data type, its scaling, its update logic, its quality flags, and the events that trigger it.

The performance envelope — update rate per point class, end-to-end latency targets, packet loss tolerance, jitter sensitivity for time-critical data.

The grid code obligations that the plant supports and the substation reports — fault ride-through behaviour, reactive power capability curves, frequency response characteristics, the events that constitute non-compliance and must be logged.

The security requirements for the link itself — IEC 62351-3 for transport-layer security on TCP/IP profiles, IEC 62351-5 for serial and derived protocols, IEC 62351-6 for the IEC 61850 protocols, certificate-based authentication, and the credential management lifecycle for the certificates and keys used at the boundary.

Time synchronisation deserves a short note because it is the one technical area where the plant and substation share an unavoidable dependency. The substation runs its own time infrastructure, typically a GNSS-disciplined master clock distributing IRIG-B over fibre or IEC 61588 (PTPv2) over Ethernet to the IEDs and the disturbance recorder. The plant runs its own time infrastructure for SCADA, historian, engineering workstations and turbine or inverter controllers. From the manufacturer's perspective, the requirement is straightforward: their equipment accepts time from an operator-provided source at the accuracy the application requires. The protocol, the source address, the path and the redundancy of that time source are operator decisions. The manufacturer states the accuracy required as a number — "synchronisation to better than 1 ms" — not as an architecture.

That is the contract at the boundary. What flows across it is bilaterally agreed and documented. What happens on either side of it is owned by the engineering discipline responsible for that side.

What the manufacturer does not get to see

The single-line diagram of the substation. The bay configurations. The protection coordination study and the relay settings. The busbar configuration and the breaker failure protection logic. The disturbance recorder configuration. The substation automation logic diagrams. The physical layout drawings of the switchyard. The list of personnel with access credentials. The substation's own cybersecurity zone and conduit diagram. The fibre infrastructure plan inside the substation. The auxiliary supply arrangement.

This is not an exhaustive list. The principle is that anything inside the substation fence — electrically or informationally — belongs to the substation engineering team and is shared on a need-to-know basis with parties whose contractual scope of work requires access to it. A wind turbine or solar inverter manufacturer's scope of work does not require it. The controller needs to ride through faults; the fault behaviour is specified at the boundary, through low-voltage ride-through curves, frequency response requirements and reactive power capability obligations expressed as parameters, without the manufacturer needing to see how those faults arise or how the substation responds to them.

There is a particular value in being clear about this with manufacturers accustomed to operating in markets where one engineering team designs the plant, the substation and the grid interface as a single integrated package. In an EU-financed project, none of those three is the manufacturer's. The plant is the asset owner's. The substation is the asset owner's substation team's, working under the connection agreement with the grid operator and under the supervision of the protection and control discipline. The grid is the grid operator's. The manufacturer is one supplier into one of those three engineering domains.

The single-line diagram example is worth lingering on because it crystallises the principle. A protection coordination study has been conducted on the assumption that the wind plant or solar plant injects current at the boundary with a defined fault behaviour. The boundary fault behaviour is what the manufacturer must deliver. The reasoning behind the protection scheme — why those particular relay settings, why that particular reclose logic, why the busbar protection is configured as it is — is the substation engineering team's working. The manufacturer does not need it to deliver the boundary behaviour, and providing it would expose engineering that is rightly held within a smaller circle of people.

At proposal stage

The principle holds in proposal documents the same way it holds in design. A bid that specifies the equipment, the boundary protocols, the boundary data points, the boundary performance, the grid code parameters supported by the equipment, and the security obligations the manufacturer will meet for the link itself — and stops there — is a bid the asset owner's substation team and the grid interface team can work with. A bid that includes a proposed single-line diagram for the substation, or that demands access to the protection coordination study for design verification, or that assumes specific relay settings, or that proposes substation cybersecurity controls beyond the boundary, is a bid that creates friction.

The friction in many early conversations is not about whether the manufacturer is competent. It is about whether they have understood that they are one engineering discipline among several, and that their authority ends at a boundary that has been deliberately drawn there. The substation team is not being unhelpful. They are being correct.

The next article moves from architectural questions to operational ones, beginning with the most consistent surprise of all: the persis-

tent VPN tunnel from the manufacturer's office to the plant, which has been the industry's default for a decade, is no longer on the table.

This article reflects the regulatory and standards landscape at publication. References to IEC 61850, IEC 62351, IEC 61588 and IEC 60870-5-104 may be superseded by revisions of those standards; national grid codes evolve continually. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

Remote access, but not the kind you remember

15 May 2026 · 10 min read · #compliance #security #industrial #oem-eu-readiness

The manufacturer's service organisation lead opens their part of the technical meeting with a question that seems uncontroversial. What is the procedure for establishing the VPN tunnel from their service centre to the plant? Will it be IPsec or SSL? Who issues the certificates? Is there a preferred vendor for the gateway appliance?

They are not asking whether there will be a tunnel. They are asking how the tunnel will be configured. The answer they expect is a routine technical exchange about gateway types, encryption parameters and credential exchange. The answer they receive is that there will not be a tunnel.

Persistent connectivity from the manufacturer's office to the plant is a 2015 architecture. It has been the industry default for a decade, embedded in long-term service agreements, condition-monitoring contracts and the operational habits of every major manufacturer in the wind, solar, battery and inverter markets. It is the way maintenance has been done. Under EU expectations, applied to an EU-financed project, it is not the way maintenance will be done.

No IPsec tunnel. No SSL VPN. No AnyDesk, TeamViewer, Splashtop, or any of the other commercial remote access tools that have proliferated in industrial maintenance. No standing connectivity of any kind, in any protocol, on any schedule. The manufacturer does not live inside the perimeter. They visit it, for as long as a specific task requires, under conditions that the asset owner controls.

This is the most operationally consequential of the architectural changes in this series, because it touches the manufacturer's busi-

ness model. The long-term service agreement presupposes the ability to observe the asset and to intervene when it underperforms. The manufacturer's first reaction is usually that without persistent connectivity, neither observation nor intervention is possible. The first reaction is wrong. Both are possible. Neither requires the manufacturer to be inside the perimeter.

Why persistent connectivity is no longer on the table

Three reasons converge.

Each persistent tunnel is a standing attack path. A VPN from a manufacturer's service network into the plant means that any compromise of the manufacturer's network — credential theft, malware on a service engineer's laptop, a successful phishing campaign against a maintenance team member, an unpatched vulnerability in the manufacturer's gateway appliance — propagates directly into the OT environment. The asset owner has no visibility into the manufacturer's internal security posture and no contractual standing to audit it deeply. Trusting that the manufacturer's network is secure enough to be inside the OT perimeter is an architectural decision the asset owner cannot validate and the lender cannot accept.

NIS2 supply chain assurance requires actively controlled access. The operator's obligations under Article 21 include managing supply chain security, controlling access to network and information systems, and being able to demonstrate to a regulator that those controls are in place. A persistent tunnel that the manufacturer can use at any time, on any schedule, without prior approval and without session-level visibility, is not actively controlled. It is access on the supplier's terms.

The audit trail is the system of record. Under NIS2 and the broader EU cybersecurity framework, the operator is expected to know who accessed what, when, for what purpose, under whose authorisation. A

persistent tunnel cannot produce that audit trail at the granularity required. It can show that a tunnel was up; it cannot reliably show which engineer did which action against which asset during which window.

The result is not a softer version of the old architecture. It is a different architecture entirely.

What replaces it

Two patterns running in parallel.

Telemetry flows out. Condition monitoring, performance analytics, fault prediction, fleet benchmarking — everything the manufacturer needs to observe the asset — runs through a unidirectional gateway from the plant out to the manufacturer's cloud. The gateway is a piece of hardware that physically permits data to flow in one direction only, by design rather than by configuration. Commercial implementations include Waterfall Security Solutions, Owl Cyber Defense and several others; the principle is the same across vendors. The manufacturer receives a continuous, near-real-time stream of operational data, processes it in their own analytics environment, runs whatever machine learning models or expert system rules they wish, and generates reports, alerts and recommendations. None of this requires inbound connectivity. The manufacturer can know everything they need to know about the asset's behaviour without ever connecting to it.

Maintenance access is brokered. When something needs to be done that requires the manufacturer to interact with the asset — investigating a fault, applying a configuration change, running a diagnostic, deploying a patch — the manufacturer requests access through a secure remote access broker operated by the asset owner. The broker is not a tunnel. It is a mediated session that opens for a defined pur-

pose, for a defined duration, against a defined target, with defined privileges, and that closes when the work is done or the time expires.

Common platforms come from two product genealogies: OT-purpose-built (Claroty xDome Secure Access, Dispel, Xage, Waterfall HERA) and IT-PAM products extended for OT use (CyberArk Privileged Session Manager for SSH and RDP, BeyondTrust Privileged Remote Access), with several others in each category. The technology differs between platforms; the operational model is the same.

What a session actually looks like

Walk through a typical maintenance session.

The manufacturer's named engineer — vetted in advance under the asset owner's onboarding process, with credentials provisioned in the asset owner's identity system — opens a service ticket. The ticket states the asset to be accessed, the work to be performed, the protocols required (SSH, HTTPS, RDP, vendor-specific tool), the duration estimated, and the change reference if the work modifies configuration.

The asset owner's authorising authority — typically a control room engineer with appropriate delegation — reviews the ticket. They confirm the asset is in a state that permits the work, that no conflicting activity is in progress, that the requested duration is reasonable, and that the change reference is approved. They authorise the session.

The engineer authenticates to the broker with multi-factor authentication. The broker establishes the session through a controlled jump host. The engineer interacts with the target through their browser or through a thin client; their own laptop never directly speaks to the target. The session is recorded in full — keystrokes, screen capture, file transfer attempts, command output.

During the session, the engineer can perform the work specified in the ticket. They cannot perform work outside that scope. The clipboard is disabled in both directions. File transfer is disabled by default and requires a separate, named justification if needed. The session has a hard expiry; if the work runs longer than expected, the engineer requests an extension and the authorising authority decides.

There is one further restriction that often surprises manufacturers. Within the session, write actions are not automatic. The conduit firewall is typically configured to permit read-only protocol behaviour by default. A manufacturer's engineer connecting to a controller's web interface may see the configuration screens correctly, but find that submitting a form returns an error — because the HTTP POST, PUT, PATCH or DELETE method has been blocked at the ICS firewall. SCADA write function codes (Modbus 5, 6, 15, 16; the IEC 60870-5-104 write ASDUs; the OPC UA Write service) are similarly blocked by default. Write privilege is requested separately, naming the specific action to be performed, with the exact target and the exact value, and is approved separately by an authority with appropriate seniority. The firewall then opens the specific write path for the specific session.

This is what IEC 62443-3-3 SR 5.1 — information flow enforcement — looks like in practice. Reads do not become writes by default. Writes are explicit, approved, time-bounded, and recorded. The principle is the same as the access principle for the session itself: nothing is implicit, nothing is durable, nothing is reusable without re-authorisation.

How availability guarantees still work

The manufacturer's concern, often expressed quietly at this point in the meeting, is that the long-term service agreement's performance commitments — availability above 97 per cent, mean time to recover within a stated number of hours, response times for critical faults —

cannot be met without persistent access. The concern is understandable but mistaken.

What persistent access actually provides for a service agreement is two things: continuous observation of the asset, and the ability to intervene quickly when intervention is required. Both can be provided without standing connectivity.

Continuous observation comes from the telemetry stream. The manufacturer's monitoring centre sees every value, every alarm, every operational state, with latency measured in seconds. Many manufacturers find that the telemetry available under a structured unidirectional architecture is richer and more reliable than what they were previously pulling through a VPN, because the gateway is engineered for high-throughput one-way flow and the data model is specified rather than improvised at each site.

Rapid intervention is a process question, not an architecture question. The mean time to first technical engagement under a brokered access model depends on how the on-call rotation, the authorisation workflow, and the broker provisioning are designed. With pre-approved emergency-response procedures, named on-call engineers whose credentials are already provisioned in the asset owner's identity system, and an authorising authority on duty round the clock, first technical engagement within ten to fifteen minutes of an alert is achievable. For service agreements written against well-designed availability targets, this is comfortably within the required response time.

The manufacturers who push back hardest on this model are often the ones whose previous service economics depended on routine remote intervention rather than scheduled, structured engagement. The shift to brokered access tends to surface a different operational rhythm — fewer ad-hoc connections, more deliberate work, more documented changes — that is, in the long run, better for the asset and

better for the audit trail. The service agreement price may need to reflect that shift. The availability guarantee does not need to.

At proposal stage

The manufacturer who proposes a persistent VPN for service access is proposing a model that the lender will not accept and the operator cannot offer. The proposal will be returned with comments asking the manufacturer to restate their service model around brokered access and unidirectional telemetry.

A bid that arrives with the new model already understood — that lists named on-call engineers who would be onboarded into the asset owner's identity system, that specifies the protocols required for typical maintenance tasks, that proposes a telemetry data model for the unidirectional flow, that estimates duty cycle for brokered sessions — is a bid that has done its homework. The conversation about how to deliver the service agreement against the asset owner's access model becomes a structured operational discussion rather than a renegotiation of the architectural premise.

The deeper point is that the manufacturer's presence inside the perimeter is not a precondition for the manufacturer's value. The value sits in the engineering judgement of the service organisation — knowing what to look for in the data, knowing what to do, knowing what risks the asset is carrying. None of that requires the engineer to be inside the network. It requires them to see the data, to communicate with the asset owner, and to act through the controlled mechanism the asset owner provides.

The next article moves to the controller board itself, and to the one component on it that asset owners have come to insist on removing physically rather than disabling in firmware: the cellular modem that arrives for "emergency support".

This article reflects the regulatory and standards landscape at publication. References to IEC 62443 may be superseded by revisions of that standard; NIS2 transposition continues to evolve across member states; named commercial products are illustrative rather than endorsements. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

Remove the cellular modem. Not disable — remove

15 May 2026 · 8 min read · #compliance #security #industrial #oem-eu-readiness

Factory Acceptance Test, day three. The asset owner's commissioning engineer is walking around an open cabinet containing the wind turbine main controller, comparing what is in front of them with the manufacturer's documentation. The cabinet is clean, well-laid out, professionally built. The documentation is complete and accurate.

Then the engineer notices something the documentation does not mention. A small SMA connector on one edge of the main control board, and, threaded behind a cable bundle, a coaxial pigtail leading to a small antenna mounted on the inside of the cabinet door.

They follow the pigtail back. A daughter card plugged into a mezzanine connector on the main board, carrying an LTE modem module of a type commonly used in industrial routers.

The manufacturer's engineer is called over. They confirm the modem is present. They confirm it is disabled in firmware. They offer to demonstrate that no traffic flows through it under normal conditions. They explain that the modem is fitted "for emergency support, in case site connectivity is unavailable during commissioning".

The asset owner's engineer makes a note in the FAT report. The note is brief: the modem must be physically removed. Not disabled. Removed.

The principle is short. Hardware that is present is a capability that exists. Hardware that is absent is a capability that does not. The two states are not interchangeable, and the difference between them is the difference between a controlled architecture and an uncontrolled one. Out-of-band channels — cellular modems, Bluetooth interfaces,

Wi-Fi adapters, hidden USB management ports, undocumented serial consoles, baseboard management controllers, NFC tags, proprietary radio links — are the single most common architectural surprise in OT acceptance testing. The remedy is not configuration. It is removal.

Why "disabled" is not enough

A disabled component can be re-enabled. Firmware is not immutable. A future firmware update — including one applied through the legitimate patch process — may re-enable the modem, either intentionally because a feature has been reinstated, or unintentionally because a regression has slipped past testing. The manufacturer's commitment that no future firmware will re-enable a particular capability is not enforceable across firmware versions, vendor mergers, changes of product strategy, or human error in a release branch. Five years into a twenty-five-year asset life, the people who made that commitment are no longer at the company, and the commitment is no longer in the change log.

Disabled does not mean inert. A cellular modem chip on the board, even with its firmware disabling the radio, still has physical RF capability if its antenna remains connected. It still has power and clock signals. It still appears in the vulnerability landscape — CVEs targeting that modem chipset still apply to the device that ships with it, whether the radio is in use or not. If the device has a SIM, the SIM may still respond to certain commands. A maintenance technician, a turbine climber, or any contractor with physical access in a cabinet has more options for reactivating a disabled radio than for installing one that was never there.

Verification is not feasible. The asset owner cannot prove a negative through inspection of a black-box device. They cannot verify that firmware actually disables the modem under all conditions and across all maintenance modes. They cannot verify that no specific UART

command, no specific GPIO sequence, no specific debug build, no specific recovery procedure will reactivate the radio. The only verifiable state is physical absence.

This is the defence-in-depth principle applied at the hardware layer: remove the capability that is not needed, rather than configure it suppressed. The capability that is not on the board cannot be enabled, cannot be exploited, cannot be re-enabled by a future firmware version, cannot appear in a future CVE, and cannot be reactivated by an attacker with physical access. The capability that is on the board, regardless of configuration, can be all of these things.

The components that arrive uninvited

The inventory is consistent across the industry.

Cellular radios — often 4G or LTE, sometimes 5G — frequently described in datasheets as "optional" or "for service support". Bluetooth, frequently positioned as a diagnostic or configuration interface, increasingly paired with a manufacturer's mobile application. Wi-Fi, sometimes for service engineer convenience, sometimes for asset tracking, sometimes for the manufacturer's cloud connectivity. Hidden USB management ports on managed switches, RTUs and HMIs, labelled "service only", typically exposing console access that bypasses every network control. Debug serial consoles on PCBs — UART headers exposed on the board, sometimes documented, sometimes not, sometimes with default credentials. Baseboard management controllers and out-of-band management interfaces on industrial PCs and SCADA servers, running their own embedded operating systems with their own attack surfaces. Proprietary radio links to weather stations, condition monitoring sensors, or vendor "ecosystem" devices. NFC tags for configuration by phone. ZigBee or LoRaWAN for low-power sensor meshes.

Two further categories deserve specific mention.

Mezzanine connectors and expansion slots are themselves a problem, separate from any card that may populate them. A controller delivered without a cellular modem but with the mezzanine connector intact, the antenna routing fitted and the power rail provisioned is a controller into which a cellular modem can be installed in fifteen minutes by anyone with physical access. A managed switch with an empty SFP cage that the asset owner does not need is a switch with an installation path for an unexpected uplink. The procurement specification should exclude the slot, not only the card.

"Out-of-band" is a term that needs careful handling. In network engineering, an out-of-band management network is generally a good thing — a separate, controlled path for managing infrastructure that survives failure of the production network. In OT security, "out-of-band" describes anything that bypasses the controlled access architecture, and is generally a bad thing. The same word, used in two communities, points in opposite directions. When a manufacturer's datasheet describes a feature as "out-of-band management" or "out-of-band diagnostic", the asset owner reads that as a defect rather than a benefit, and asks for it to be removed.

Where this conversation belongs

Three stages, three very different costs.

At specification stage, the conversation costs almost nothing. The procurement specification explicitly excludes cellular, Bluetooth, Wi-Fi, NFC and proprietary radio. It excludes hidden USB management interfaces. It excludes optional expansion slots that could host such components later. It requires the manufacturer to declare every wireless capability — present, absent, populated, unpopulated — in the compliance documentation. The manufacturer's product variant for the project is configured accordingly at the factory, and the variant appears in the bill of materials.

At factory acceptance test, the conversation gets expensive. The asset owner's commissioning engineer physically inspects every product variant. A spectrum analyser is used to detect active radios in the laboratory. PCBs are inspected for daughter cards, populated chips, RF traces and antenna connections. Findings result in a non-conformance, the equipment is returned to the manufacturing line, and the schedule slips while modifications are made. The cost is measured in weeks of programme delay and sometimes in re-certification of the modified equipment under the manufacturer's quality system.

At site acceptance test, the conversation is the most expensive. Equipment is on site. Sometimes there are hundreds of units installed. Retrofit means either returning units to the factory at significant logistics cost, or sending factory engineers to the site to perform modifications under field conditions, with all the quality control implications that field rework brings. Programme impact is severe; commercial impact can be severe enough to threaten financing milestones.

The pattern is consistent across many of the topics in this series, but it is starkest here. Specification is cheap, FAT is costly, SAT is painful. The discipline of moving the conversation to the cheapest stage is, in many ways, the entire procurement learning for non-EU manufacturers selling into EU-financed projects.

At proposal stage

A manufacturer's bid that arrives with a wireless-and-out-of-band capability matrix already attached — declaring every chip on the board with any RF capability, every connector that could host one, every management interface, every debug port, with a clear statement of which are populated, which are not, which can be removed for the project variant, and which cannot — is a bid that has saved everyone several months. The conversation that follows is whether the manu-

facturer's standard variant can be supplied or whether a project-specific build is needed, and if the latter, what the schedule and cost implications are.

A manufacturer's bid that does not mention any of this, on the basis that the equipment "supports" remote management or diagnostic access, is a bid that signals an architectural assumption — that the manufacturer's idea of a complete delivery includes capabilities the asset owner will need to subtract. That subtraction will happen at the worst possible point in the programme.

The principle is one the manufacturer's own product engineering team will recognise on reflection. Every chip on the board is a maintenance liability, a CVE exposure, a power draw, a thermal load, a regulatory burden, a supply chain dependency. Removing the ones that the deployment does not need is not a concession. It is sound engineering. The asset owner asking for the modem to be removed is asking for the same kind of discipline the manufacturer's own value engineering team would apply for a different reason.

The next article picks up a topic the manufacturer's procurement team will recognise immediately, even if their security organisation has not yet brought it to them: the IEC 62443-3-2 zone and conduit risk assessment that the L0/L1 system integrator is expected to produce for the project, and what it actually contains.

This article reflects the engineering and procurement practices common in EU-financed renewable energy projects at publication. Specific examples of components and interfaces are illustrative; the principle of physical removal over firmware disablement applies regardless of which particular technologies appear in any given product. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

What 'system integrator for L0/L1' actually means under 62443

15 May 2026 · 12 min read · #compliance #security #industrial #oem-eu-readiness

A design coordination meeting between the manufacturer's L0/L1 engineering team and the asset owner's L2-and-above architecture team. The asset owner's architect explains the immediate question: in order to specify the conduit firewall between the manufacturer's process zone and the upstream plant network, the architect needs the manufacturer's zone and conduit risk assessment. Without that document, the conduit specification has nothing to anchor to — no defined Target Security Level on the manufacturer's side of the conduit, no documented risk that the firewall is mitigating, no mapping between the system requirements in IEC 62443-3-3 and the actual configuration that needs to be implemented.

The manufacturer's team confers briefly. They ask which document, specifically, is being requested. They are familiar with the IEC 62443 series in the way most engineering teams are familiar with standards they have heard cited but never authored — they know the number, they know it concerns industrial cybersecurity, they have not previously been asked to produce documents under it.

The asset owner's architect explains. The manufacturer's team takes notes. The meeting ends with an action item: the manufacturer will revert with a plan for producing the requested deliverables. Two weeks later, the response comes through procurement. The manufacturer would like to understand whether the deliverables can be supplied as a separate, costed scope item, on the basis that they are not part of the equipment supply agreement as currently drafted.

This is not an unusual conversation. It is the conversation. The asset owner's L2-and-above design assumes that the L0/L1 integrator has produced specific cybersecurity documentation, because IEC 62443-3-2 — the standard that governs system risk assessment and zone-and-conduit design — explicitly places the responsibility for that documentation on the party performing the integration. When the equipment manufacturer also acts as the L0/L1 integrator, which is normal for wind turbine, solar inverter, battery and hybrid plant suppliers, the integrator obligations flow to them. They do not go away because the standard is unfamiliar.

The two standards that bear on this role

Two parts of the 62443 series are operationally relevant.

IEC 62443-2-4 is the security programme requirements for IACS service providers. It defines what a system integrator's own security management programme must contain — staffing, awareness, change management, patch management, audit log management, several other process areas — and what evidence is expected to demonstrate conformance. A manufacturer acting as integrator is, for the purposes of this standard, a service provider, and is expected to maintain a security programme aligned to it.

IEC 62443-3-2 is the security risk assessment and system design standard. It defines the process by which the integrator analyses the system under consideration (SuC), partitions it into zones and conduits, conducts a risk assessment, assigns Target Security Levels to each zone, and produces the documented outputs that downstream engineering depends on. This is the standard that produces the deliverables the asset owner needs.

Two further parts of the series sit alongside these and are referenced rather than authored.

IEC 62443-3-3 specifies the system requirements organised by foundational requirement (FR1 through FR7) and security level (SL 1 through 4). The integrator does not author 62443-3-3 — it is a published standard — but they apply it. The Target Security Levels assigned in the 3-2 risk assessment become demands on the system, which are then mapped to the 3-3 system requirements that must be met at each level.

IEC 62443-4-2 specifies the technical security requirements for the components themselves, organised the same way as 3-3 but addressed to the component supplier rather than the system. A component's Capability Security Level — SL-C — is a property of the component, certified by independent test where possible. The components a manufacturer delivers into the project must have SL-C values at least equal to the Target Security Level (SL-T) of the zone they sit in. Where they do not, compensating controls must be designed and documented.

The relationship between these four parts is the architecture of the entire conversation. 2-4 says what the integrator must do as an organisation. 3-2 produces the risk assessment and the zone-and-conduit design. 3-3 defines the system requirements at each Security Level. 4-2 defines the component capabilities that allow the system to meet them. The L0/L1 integrator sits in the middle of all four.

The documents an L0/L1 integrator produces

The 3-2 process produces a series of artefacts the asset owner expects to receive and review. The IEC 62443 evidence pack post lists the asset-owner-side counterparts; the integrator-side artefacts below feed those packs.

The System under Consideration definition. A scoping document that states clearly what is included in the integrator's L0/L1 system — which assets, which interfaces, which protocols, which physical and

logical boundaries — and what sits outside, particularly the boundary at which the integrator's responsibility ends and the asset owner's begins. The SuC is the foundation; if it is wrong, everything downstream is wrong.

The high-level risk assessment. A first-pass analysis determining whether the SuC presents sufficient risk to warrant detailed assessment. For an L0/L1 system in a utility-scale renewable energy plant, the answer is always yes, but the high-level assessment documents the reasoning and supports the initial zoning decisions.

The initial system partitioning. The first cut at zones and conduits, identified from the SuC. Zones group assets that share a common security level and exposure profile; conduits are the controlled communication paths between zones. For a typical wind turbine, the zoning might include the per-turbine controller zone, the turbine cluster network zone, the plant SCADA zone, and the boundary to the upstream IDMZ — with conduits between each. For a solar plant, similar with inverters in place of turbine controllers; for a battery system, similar with the BMS in the role of the controller.

The detailed risk assessment per zone and conduit. For each zone, an analysis of the threat scenarios that could affect it, the likelihood of those scenarios, the consequences (operational, safety, financial, environmental), and the residual risk after the planned controls. For each conduit, an analysis of the communication that crosses it, the threats specific to that communication path, and the controls applied to mitigate them.

The Target Security Level assignment. For each zone, an SL-T value of 1, 2, 3 or 4, applied independently to each of the seven foundational requirements. A turbine controller zone in a utility-scale plant might be assigned SL-T 3 for system integrity (FR3) and timely response to events (FR6), SL-T 2 for use control (FR2) and restricted data flow (FR5), and lower levels for the remainder. The SL-T values

quantify what the zone must be able to defend against and direct the choice of components and configurations.

The Cybersecurity Requirements Specification. The output document that captures all of the above and serves as the input to detailed design. It states the SuC, the zones, the conduits, the SL-T per zone per foundational requirement, the threats considered, the controls required, and the residual risk accepted by the asset owner. It is the document that the asset owner's L2 design references when specifying conduit firewalls, IDMZ rules, monitoring requirements and incident response procedures.

The SL-C mapping. For each component the integrator delivers into the SuC, a statement of the Capability Security Level it provides per foundational requirement, with evidence — typically a certification or evaluation report — supporting the claim. Where the SL-C is lower than the SL-T of the zone, compensating controls are documented in the Cybersecurity Requirements Specification and the residual risk is formally addressed.

These are not optional artefacts. They are the design basis for everything downstream. Without them, the asset owner cannot specify the L2-and-above architecture, the lender cannot evidence cybersecurity due diligence, and the project cannot demonstrate 62443 conformance to any independent assessor.

What Security Levels actually mean

The four-level scale in 62443 is not a generic risk rating. It corresponds to specific threat capabilities, defined in 62443-1-1 and applied consistently across the series.

Security Level 1 is protection against casual or coincidental violation. Default credentials changed, basic access control, basic logging. The floor of any responsible deployment.

Security Level 2 is protection against intentional violation using simple means with low resources, generic skills and low motivation. The threat actor in scope is an opportunistic insider or external party using widely available tools.

Security Level 3 is protection against intentional violation using sophisticated means with moderate resources, IACS-specific skills and moderate motivation. The threat actor is a serious adversary with industry-specific knowledge — disgruntled former employees with privileged knowledge, organised criminal groups targeting industrial control systems, regional state-affiliated groups.

Security Level 4 is protection against intentional violation using sophisticated means with extended resources, IACS-specific skills and high motivation. The threat actor is a peer state intelligence service, an advanced persistent threat with multi-year campaigns, organised crime operating at nation-state scale.

For utility-scale renewable energy plants in EU-financed projects, particularly in regions of geopolitical exposure, SL-T 3 is the typical floor for the foundational requirements that bear on system integrity, restricted data flow and timely response to events. SL-T 2 may be acceptable for some lower-criticality zones. SL-T 4 is rarely required outside specific national infrastructure designations or against named threat actors.

The translation matters because it determines what the integrator must deliver and what the asset owner must build around it. SL-T 3 components have particular capability requirements — strong identity and authentication, integrity-protected communication, comprehensive audit logging, hardened configurations, support for centralised key management — that are simply not present in components designed for less demanding markets. A manufacturer whose products are well-suited to an SL-T 2 deployment in a domestic market may need different components, or substantial additional engi-

neering and compensating controls, to deliver into an SL-T 3 deployment in an EU-financed project.

When and how this work happens

The 62443-3-2 process is not a final-acceptance deliverable. It is a design basis, which means it must exist before downstream design depends on it.

The typical sequence runs as follows.

Pre-final-investment-decision, during bid and conceptual design: an initial scoping exercise. The System under Consideration is sketched, the high-level risk assessment is conducted, the initial zoning is proposed. This is sufficient for the lender's early due diligence and for the asset owner's L2 conceptual design.

Post-final-investment-decision, in early detailed engineering: the detailed risk assessment is conducted, the zones and conduits are finalised, the SL-T values are assigned and signed off. The Cybersecurity Requirements Specification reaches its first formal version. The asset owner's L2 design begins to firm up around this baseline.

Through detailed engineering and procurement: the SL-C evidence is assembled for each component, the Cybersecurity Requirements Specification is updated as design decisions are made, the conduit specifications are finalised. The L2 firewall rule base, IDMZ configuration and monitoring infrastructure are designed against the document.

At factory acceptance test: the components are verified against the SL-C claims and the Cybersecurity Requirements Specification expectations. Findings result in non-conformances or accepted residual risks; either way they are documented.

At site acceptance and commissioning: the integrated system is verified against the Cybersecurity Requirements Specification, the opera-

tional controls are demonstrated, the residual risks are formally accepted by the asset owner.

Typical duration from kick-off to a signed-off baseline Cybersecurity Requirements Specification is six to ten weeks of focused work for a single utility-scale plant. Less if the manufacturer has produced 62443-3-2 documentation for previous projects and has templates and prior examples to start from. Considerably more if the manufacturer has never produced one and is starting from a standing position.

Review parties typically include the asset owner's security architect (the primary internal reviewer), the asset owner's engineering and operations teams (for the operational consequences), the project's IEC 62443-conformant cybersecurity consultant if one is engaged, the lender's technical and cybersecurity adviser, and where independent assurance is required, a third-party certification body. ISASecure CSA is the most common scheme path to SL-C evidence; assessment services for the 62443 series are offered by TÜV SÜD and exida among others, with DNV and Bureau Veritas also active in specific market segments.

At proposal stage

A manufacturer's bid that arrives with the 62443 deliverables already accounted for — that names the integrator role explicitly, that proposes a 3-2 work plan with timeline and review gates, that lists the SL-C evidence for the components being offered, that identifies any gaps between component SL-C and likely zone SL-T and proposes how those gaps will be closed — is a bid that demonstrates the manufacturer has done this before, or at least knows what doing it looks like. The conversation that follows is about scope, schedule and resourcing, not about whether the work is required.

A manufacturer's bid that does not mention 62443 at all, or that mentions it as a future deliverable to be scoped separately, signals one of

two things. Either the manufacturer is not yet equipped to take on the L0/L1 integrator role under EU-financed terms, or they intend to take it on but have not yet recognised the work that role entails. Both are surmountable. The first is the work of months, beginning with a 62443-2-4 service-provider programme; the second is the work of weeks, beginning with a project-specific 3-2 work plan and an accredited consultant alongside the engineering team. But both must be addressed before contract signature, because nothing in the L2-and-above design proceeds without the 3-2 deliverables, and the project schedule does not pause to wait for them.

The recurring observation throughout this series is that EU cybersecurity expectations are not as forbidding as they sometimes appear in the first conversation. The 62443 documentation is the clearest case. The process is well-defined. The deliverables are clear. The standards exist. Consultants accredited to assist exist. The work, once understood, is routine engineering of a kind the manufacturer's organisation already does for other purposes — risk assessment, system architecture, requirements traceability — under a different name and a different framing.

The next article moves to a topic the manufacturer's organisation may find more cultural than technical: the public-facing vulnerability disclosure programme that the Cyber Resilience Act will require, and why "email us if you find a problem" is not a programme.

This article reflects the IEC 62443 series at publication. The standard continues to evolve; in particular 62443-3-2 and 62443-4-2 have been subject to revision discussion through the IEC technical committee. References to commercial certification bodies are illustrative rather than endorsements. Specific arrangements should be reviewed by

qualified legal counsel rather than against this article. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

Email is not a vulnerability disclosure programme

15 May 2026 · 11 min read · #compliance #security #industrial #oem-eu-readiness

A late-stage proposal review. The asset owner's security architect asks the manufacturer how vulnerabilities affecting the deployed equipment will be communicated during the operational life of the asset. The manufacturer's response is helpful and concrete. They will maintain a distribution list of customer security contacts. When a vulnerability is identified, the security contacts will be notified by email. The notification will include the affected product, the severity, and the patch when one is available.

The architect notes that this is useful, but does not address the requirement. The Cyber Resilience Act requires a coordinated vulnerability disclosure programme accessible to the public — not only to existing customers. The manufacturer's team is briefly puzzled. Why does the public need access? The vulnerabilities affect the asset owner's equipment. Surely the asset owner is the appropriate audience?

The architect explains. Vulnerability disclosure is not a customer-relationship function. It is a public function, governed by international standards, expected of every serious industrial equipment manufacturer, and required in operational form by EU regulation from September 2026. The asset owner is one beneficiary of the programme. The security research community, downstream asset owners on the second-hand market, parallel deployments at other operators, the national CERT community, sector ISACs, vulnerability scanner vendors, insurance underwriters, regulators — all are beneficiaries. Restricting disclosure to a private distribution list serves none of them.

A vulnerability is information. The asset owner's interest is in receiving it; the wider ecosystem's interest is in being able to act on it; the security research community's interest is in having a channel through which to report what they find. None of these interests is served by private email. Each requires a public, persistent, structured mechanism for vulnerability disclosure — a programme rather than an inbox.

Why disclosure is a public function

The Cyber Resilience Act, in Article 13 (via Annex I Part II point 5), requires manufacturers of products with digital elements to operate a coordinated vulnerability disclosure policy. The phrase "coordinated vulnerability disclosure" carries specific meaning under ISO/IEC 29147 — the international standard that defines what such a policy must contain — and the operational counterpart, ISO/IEC 30111 , which defines what a manufacturer's internal vulnerability handling process must look like. Together these two standards set out the architecture of a credible disclosure programme.

What the standards require, and what the Cyber Resilience Act will enforce from December 2027, is structural rather than discretionary. The manufacturer publishes a clear policy stating how vulnerabilities can be reported, by whom, under what conditions. The manufacturer provides a public channel for receiving reports, accessible to anyone who finds a vulnerability without requiring an existing customer relationship. The manufacturer commits to a process for triaging, investigating, remediating and publishing information about vulnerabilities, with stated timelines. The manufacturer publishes advisories once remediations are available, in a format that downstream parties can monitor and act on.

Three reasons make the public posture non-negotiable.

Security researchers do not have customer relationships. A researcher in Helsinki, a graduate student in Tel Aviv, a security team at an unrelated company, an independent bug bounty hunter — none of them are customers of the manufacturer, and none should need to be. They have found a vulnerability and they need a way to tell the manufacturer about it without going through a sales call. Private email distribution lists serve customers. Researchers reach manufacturers through the public channel or, if no public channel exists, through public disclosure on a fixed timeline. The manufacturer's choice is between operating the channel themselves or having the disclosure happen without them.

Downstream parties depend on advisories. A wind turbine installed at the project in question is one of perhaps several hundred globally of the same model. A vulnerability affecting it affects them all. The other asset owners — whether they bought directly from the manufacturer, inherited the equipment through acquisition, or operate it under second-life arrangements — need access to the same advisory. National CERTs and sector Information Sharing and Analysis Centres feed off published advisories to provide guidance to their constituencies. Insurance underwriters cross-reference advisories against the equipment they cover. Vulnerability scanning tools ingest advisories into their detection databases. A vulnerability not published is not visible to any of this infrastructure.

Equivalence with peers. The major industrial control manufacturers — Siemens, ABB, Schneider Electric, Rockwell, GE Vernova, Hitachi Energy, Mitsubishi, Yokogawa, several others — all operate published vulnerability disclosure programmes with security.txt files, advisory pages, CVE issuance, PGP keys, and acknowledged researcher credits. A manufacturer that lacks these is signalling either that they have not yet reached the operational maturity of peers, or that their internal posture toward vulnerabilities is more defensive than collabora-

tive. Both readings are damaging at proposal evaluation, and both are increasingly difficult to defend at lender due diligence.

What a PSIRT actually does

A Product Security Incident Response Team — PSIRT, to use the common acronym — is the organisational function that handles vulnerabilities affecting the manufacturer's products. The work has six phases.

Intake. A vulnerability arrives through the public channel — a web form, an email address, an encrypted message. The intake process acknowledges receipt within a stated timeline (typically three business days), records the report in a tracking system, and assigns initial triage.

Triage. An analyst confirms the report is genuine, identifies which product or products are affected, estimates initial severity, and decides whether to escalate. In the operational experience reported by established PSIRTs, a substantial fraction of reports turn out to be duplicates, configuration issues or scope mismatches; the residue requiring engineering action varies by product.

Coordination. The PSIRT works with the reporter on a disclosure timeline. Standard practice under ISO/IEC 29147 is 90 days from initial report to public disclosure, with extensions negotiable if remediation is complex. The reporter and the manufacturer agree on what will be published and when; if they cannot agree, the standard provides for the reporter to disclose unilaterally after the standard timeline. A PSIRT that fails to engage substantively within the 90 days finds itself responding to public disclosure rather than coordinating it.

Remediation. Engineering develops the fix. This is the most variable phase — a configuration change might take a day, a firmware update with full regression testing might take months. The PSIRT tracks

progress and updates the reporter on timeline, while parallel work prepares the advisory text, the customer notifications and the patch delivery infrastructure.

Disclosure. Once the remediation is available, or once the disclosure deadline passes, the PSIRT publishes the advisory. The advisory contains the affected products, the affected versions, a CVE identifier, a CVSS score, a description of the vulnerability written for the technical audience that needs to act on it, the remediation guidance, and any workarounds for asset owners who cannot patch immediately.

CVE issuance. The Common Vulnerabilities and Exposures programme, operated by MITRE under sponsorship from the US Cybersecurity and Infrastructure Security Agency, assigns unique identifiers to publicly disclosed vulnerabilities. A manufacturer that operates as a CVE Numbering Authority — a CNA — can assign CVE identifiers to their own products' vulnerabilities, which signals operational maturity and integrates the manufacturer into the global vulnerability tracking ecosystem. Applying to become a CNA is a process MITRE manages openly. For industrial equipment manufacturers selling into EU-financed projects, it is the expected end state.

The minimum-viable PSIRT

The full operational programme is well-defined. The minimum-viable initial deployment is achievable in four to six weeks of focused work, most of which is editorial and process design rather than engineering.

A security.txt file at the root of the manufacturer's primary domain — at `/.well-known/security.txt` under RFC 9116 . The file declares the contact for vulnerability reports, the encryption key for confidential reports, the policy URL, the acknowledgments URL where credited researchers are listed, and the preferred languages for communication. A dozen lines of plain text; one of the simplest deliverables in the en-

tire programme; the file that researchers actively look for as their first signal that a manufacturer takes disclosure seriously.

A `/security/advisories` page on the manufacturer's website. A clean, navigable list of published advisories, each with a unique identifier, an affected products list, a CVE reference where applicable, a CVSS score, the publication date, the remediation status, and links to firmware downloads or configuration guidance. Static HTML is sufficient; the page does not need to be application-driven. The publication discipline matters more than the technology.

A PGP key, or modern equivalent, for encrypted communication. Researchers reporting sensitive vulnerabilities require a confidential channel. A published PGP key with a known fingerprint, accessible from the security.txt file, and a separate encrypted-mailbox infrastructure for receiving reports, addresses this requirement. Age keys are increasingly accepted as a modern alternative, though PGP remains the default in security research communities.

A coordinated vulnerability disclosure policy document, aligned to ISO/IEC 29147. The policy states the manufacturer's commitments to researchers — that good-faith reports will not result in legal action, that disclosure timelines are stated and adhered to, that credit will be offered for valid reports unless the reporter prefers anonymity. Without these commitments, serious researchers will not engage. With them, the channel becomes productive.

An RSS or Atom feed of advisories. Downstream consumers — CERTs, ISACs, vulnerability scanning vendors, automated procurement tooling, insurance risk engines — ingest advisories through feeds rather than human visits. A manually maintained advisory page that requires browser visits will not reach the infrastructure that needs it.

A CVE Numbering Authority application, in progress or completed. Manufacturers that operate as CNAs can assign their own CVE identi-

fiers, which speeds the disclosure process, demonstrates a degree of maturity that procurement teams notice, and embeds the manufacturer into the wider vulnerability tracking ecosystem.

These six artefacts, taken together, constitute the minimum-viable PSIRT. They are not the full operational programme — that requires staffing, internal procedures, executive ownership, engineering escalation paths, integration with the development lifecycle, and continuous engagement with the security research community. But they are the visible artefacts the asset owner, the lender, and the security researcher will look for. A manufacturer that has them, even if the operational programme behind them is still maturing, has crossed the threshold of credibility. A manufacturer that has none of them has not. The IEC 62443 evidence pack post covers the asset-owner-side artefacts that consume the PSIRT's outputs.

At proposal stage

A manufacturer's bid that includes the URL of an existing security.txt and advisory page, the disclosure policy document, the CNA reference, and a brief description of the PSIRT function, is a bid that meets the Cyber Resilience Act's vulnerability handling provisions at face value. The conversation that follows is about how project-specific advisories will be handled — escalation paths to the asset owner, integration with the project's incident response procedures, contractual notification timelines that may run faster than the public disclosure timeline — not about whether the function exists.

A manufacturer's bid that offers a customer-only email distribution list, or that proposes to set up a disclosure programme after contract signature, signals a gap that procurement and the lender will both flag. The gap is not unbridgeable. The four-to-six-week timeline to a minimum-viable PSIRT is short enough that it can run in parallel with

contract negotiation. But the work needs to start before the bid is competitive, not after.

The cultural shift, finally, is the one this piece has been circling. Traditional vendor instinct treats vulnerabilities as private matters — minimise public attention, brief affected customers quietly, work issues outside view. The mature posture, now codified in the Cyber Resilience Act and embedded in the operating practice of every major industrial control manufacturer, is the opposite. Vulnerabilities are inevitable in any complex product. How a manufacturer handles them — promptly, transparently, in dialogue with the security research community, with credit for those who report in good faith — is the signal procurement teams now read. A manufacturer that publishes advisories is a manufacturer that knows what is in their product, that has the engineering capability to fix what they find, and that trusts their organisational maturity to let the information be seen.

The procurement criterion has been quiet about this until recently. It is no longer quiet.

The next article picks up the technical foundation that vulnerability disclosure rests on, and that the Cyber Resilience Act will require independently: the software bill of materials that catalogues what is actually inside the firmware shipping in every product, and why most non-EU industrial manufacturers have never produced one.

This article reflects the regulatory and standards landscape at publication. The Cyber Resilience Act's implementing acts continue to be issued through 2026–2027 and may alter specific disclosure obligations. References to ISO/IEC 29147, ISO/IEC 30111 and RFC 9116 may be superseded by revisions. Named manufacturers and certification bodies are illustrative rather than endorsements. Specific trans-

actions should be reviewed by qualified legal counsel rather than against this article. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

What is in your firmware: the bill of materials nobody asked for before

15 May 2026 · 12 min read · #compliance #security #industrial #oem-eu-readiness

The manufacturer's lead embedded engineer is running an SBOM generation tool against the firmware image of one of their main controller products. The tool — Syft, Trivy, or a commercial equivalent — works by analysing the binary contents of the firmware and identifying known components by signature, file structure and content matching. It runs for several minutes and produces a JSON file with over a thousand entries.

The engineer scrolls through. The first hundred entries are unsurprising — BusyBox, glibc, OpenSSL, Dropbear, the Linux kernel, the lwIP stack, several standard libraries the team explicitly maintains. The next several hundred are unfamiliar. ZeroMQ. mbedTLS in addition to OpenSSL. An old version of cJSON. A Bluetooth stack the engineer is certain has never been used in this product. Half a dozen libraries with German-language comments. A build of OpenJDK the team has no record of including. Several proprietary binary blobs identified only by their cryptographic hashes.

The lead engineer reads the list twice. Most of it is correct — these libraries are in the firmware. Some have been there since the original product version six years ago, inherited from a base image the team adopted without auditing. A few entries genuinely surprise: components the team thought had been removed, versions they thought had been upgraded, modules they did not know were present at all.

This is the discovery moment, and for most non-EU industrial manufacturers, it is the first time it happens. The Cyber Resilience Act will require it to happen continuously, every release, with the resulting

bill of materials made available to the asset owner and to market surveillance authorities. The CRA's SBOM obligation sits in Annex I Part II point 1, imposed via Article 13. The clause is short and specific. Most of the work it implies is not.

A software bill of materials is what its name suggests — a list of every software component present in a product, along with version numbers, license information, supplier identity, and, where available, cryptographic hashes. The format is structured and machine-readable. The two standard formats are SPDX (Software Package Data Exchange, ISO/IEC 5962) and CycloneDX (an OWASP project). Both are widely supported by tooling, by procurement processes and by the Cyber Resilience Act's implementing acts.

What the SBOM enables is not, in itself, security. It is visibility. With an SBOM, the manufacturer knows what is in their firmware. With access to an SBOM, the asset owner knows what is in the equipment they are operating. With a published SBOM, the security research community can cross-reference vulnerabilities against deployed assets. Without one, none of these parties can act on vulnerability information except by going back to the manufacturer and asking individually, every time.

What an SBOM is, and what it is not

The SBOM is an inventory. It states what is present. It does not state whether what is present is secure. It does not state whether it is current. It does not state whether the components have been patched. It does not state whether the cryptographic algorithms used are still considered fit for purpose. It does not state whether the open source licenses have been complied with. It does not state whether the supplier of any third-party component is still trading.

These are all separate analyses that consume the SBOM as input. The vulnerability scanner consumes it to identify which deployed CVEs

apply. The license compliance tool consumes it to verify obligations. The supply chain risk tool consumes it to identify dependencies on sanctioned or untrusted suppliers. The asset owner's procurement team consumes it to cross-reference against organisational policies on component preferences.

A useful framing: the SBOM is the bill of materials for a complex machine. A car has a bill of materials. It lists every part — brake pads, alternator, ignition coil, fuel pump — by part number, supplier and version. The bill does not state whether the brake pads are worn, whether the alternator is approaching end of life, or whether any specific component has been subject to a recall. It states what is there. Other systems — service records, recall databases, condition monitoring — handle the rest.

The Cyber Resilience Act does not require manufacturers to certify their components as vulnerability-free. It requires them to know what they ship, and to make that knowledge available to the parties downstream who need to act on vulnerabilities when they emerge. The SBOM is the foundational artefact of that visibility.

What goes in firmware that catches manufacturers out

For most non-EU industrial manufacturers, the first SBOM is an education. Several categories of component routinely surprise.

Open source libraries are the largest and most consistent category. A typical embedded Linux firmware image contains several hundred open source components, often with deep transitive dependencies. The build system pulled them in. The engineering team did not necessarily review them individually. Some have been there since the original product version. A first-pass SBOM frequently identifies open source components the team had no specific knowledge of including, often because they were inherited from a base image, a board support package, or a vendor-supplied software development kit.

White-label boards and reference designs are the most awkward category. The manufacturer's product may use a System-on-Module purchased from a third-party vendor — Toradex, Variscite, Compulab, several others — that ships with its own firmware, its own bootloader, its own kernel build, its own pre-installed components. The manufacturer integrated the module into their product but did not author the lower-level firmware. The SBOM must cover everything in the deployed product, including the module's contribution. Generating this part of the SBOM requires either cooperation from the module supplier (who may or may not provide one) or independent binary analysis of the module's firmware (which is technically possible but rarely standard practice).

Chipset vendor blobs are the third surprise. Cellular modems, GNSS chips, Wi-Fi chipsets, Bluetooth controllers, FPGAs, hardware accelerators — almost all of these ship with proprietary firmware blobs that load at boot and execute on the chip itself. The vendors typically provide the blob as a binary, sometimes with an opaque license. The SBOM should identify the blob, its version, its supplier and any known vulnerabilities — but the manufacturer's view into what the blob actually contains is, by design, limited.

Inherited firmware from previous corporate transactions is the fourth, less common but more difficult. A product line acquired through corporate transaction, an OEM rebadge agreement, or a long-standing technology licensing arrangement may carry firmware components for which the operating manufacturer does not have full provenance. The components are in the product. Documenting them in the SBOM is the right thing to do. Tracing their full origin may require effort the engineering team has not previously expended.

How an SBOM is actually produced

Two complementary approaches, usually run together.

Build-time SBOM generation integrates with the manufacturer's build system. As the firmware is compiled, the build system records every package, library and source file that contributes to the final image. Tools that support this include Yocto's built-in SBOM generation, SPDX support in Buildroot, and the Software Composition Analysis tools that integrate with continuous integration pipelines — Snyk, Sonatype, Black Duck, Mend, several others — alongside the open source CycloneDX tooling. Build-time generation is more accurate because it knows what was intentionally included. It requires the manufacturer to have a build system they can instrument — which is normal practice for mature engineering organisations, but not universal.

Binary analysis examines a compiled firmware image and identifies components by signature, file structure and content matching. Tools include Syft, Trivy, Binwalk in combination with other identification tools, and commercial binary analysis platforms (ReversingLabs, Cybellum, Finite State, several others). Binary analysis is less accurate at the version level but more honest about what is actually in the deployed binary — it catches components that the build system may not be aware of, including those introduced by third-party modules and chipset blobs. Many manufacturers run both approaches and reconcile the results, treating the union as the authoritative SBOM.

The first generation is the hardest. The team finds components they did not know were present, discovers version mismatches between what was intended and what shipped, identifies licensing situations that had not been raised. The second and subsequent generations are routine, because the engineering process and the build system are now instrumented to produce SBOMs as a continuous output rather than as a one-off exercise.

The Cyber Resilience Act expects SBOMs to be available for every product release. The expectation is that SBOM generation is a normal

output of the development lifecycle, not a separate exercise scheduled before regulatory submission.

Who reads it, and why

The lender's risk team often reads the SBOM before the asset owner's security team does. Three reasons.

First, the lender's risk team is conducting cybersecurity due diligence at a stage when the asset owner's security team is still being assembled for the project. The SBOM is one of the earliest artefacts that gives the lender a concrete view of what they are financing. A clean SBOM that lists current versions of well-maintained components, with no unsupported dependencies and no components from sanctioned suppliers, is a strong early signal of engineering discipline.

Second, the lender's risk team cross-references the SBOM against EU and US sanctions lists, dual-use export control schedules under EU Regulation 2021/821, and the supplier lists maintained by their internal compliance function. Components from sanctioned suppliers, or from suppliers whose beneficial ownership is unclear, create compliance exposure for the lender. The SBOM reveals supplier identity at a granularity that contractual due diligence does not.

Third, the lender's risk team uses the SBOM to assess long-tail risk. A firmware that depends on a library last updated in 2017, maintained by an individual who no longer responds to issues, is a risk that compounds over the twenty-five-year asset life. The SBOM surfaces this kind of dependency in a way that no other artefact does.

The asset owner's security team consumes the SBOM differently. They feed it into their vulnerability management infrastructure, cross-reference it against advisories from the manufacturer's PSIRT and from independent sources, monitor for new CVEs against the listed versions, and prioritise patching based on actual exposure rather than generic severity scores. With an SBOM, vulnerability manage-

ment is a continuous process. Without one, it is a series of after-the-fact crises. The evidence pack post covers the asset-owner-side artefact set this consumption feeds into.

A specific artefact that pairs with the SBOM is the Vulnerability Exploitability eXchange document — VEX — which states, for each known CVE against a listed component, whether the vulnerability is actually exploitable in the manufacturer's product. A library may be present in the firmware but its vulnerable function may not be reachable from any code path the product actually uses. The VEX statement says so, with reasoning. Without VEX, the asset owner sees every CVE against every component and must assume the worst. With VEX, the asset owner can prioritise patching against actual exposure. The CRA's implementing acts will reference VEX or an equivalent mechanism as part of the manufacturer's vulnerability handling obligation.

The security research community consumes the SBOM to focus their effort. A researcher who finds a vulnerability in a specific library, then cross-references published SBOMs to identify which products are affected, can coordinate disclosure efficiently. Without published SBOMs, the same researcher has to test every product they suspect — far less efficient, far less likely to result in coordinated disclosure across the affected fleet.

A final consumer worth mentioning is the manufacturer's own engineering team. The discipline of producing an SBOM continuously, and reviewing it across releases, surfaces internal issues the team would otherwise not see — supply chain creep, unintended dependencies, components that have not been updated, license obligations that have accumulated. The SBOM is, in this sense, also an internal management tool.

At proposal stage

A manufacturer's bid that includes a sample SBOM for the relevant product, or a commitment to deliver SBOMs for every product release in SPDX or CycloneDX format with VEX statements for material CVEs, is a bid that has resolved a topic the lender's technical adviser will otherwise raise as an outstanding item. The conversation that follows is about delivery cadence, format preference, storage location and update mechanism — not about whether the SBOM will be produced.

A manufacturer's bid that does not mention SBOMs, or that proposes to provide a high-level inventory document at first article inspection, signals one of two things. Either the manufacturer's build system does not produce SBOMs as a normal output and the team is hoping the question does not arise, or the manufacturer is uncertain what an SBOM is and is hoping the topic does not become a procurement criterion. Both are surmountable. The first is the work of weeks, integrating SBOM tooling into an existing build system. The second is the work of one focused engineering review, after which the team will produce SBOMs routinely.

There is a useful framing for the manufacturer's internal conversation. Every chip on the board is, in some sense, a supplier. Every library in the firmware is a supplier. Every binary blob from a chipset vendor is a supplier. The product manager already knows their physical supply chain — who delivers what part, in what volume, with what lead time, at what cost, against what quality history. The SBOM is the equivalent map for the software supply chain. Mature engineering organisations have been building this map for years, under various names. The Cyber Resilience Act simply requires it to be visible to others.

This is the eighth article in the series, and the last one that asks the manufacturer to produce a new artefact. The pieces that follow de-

scribe disciplines the manufacturer's organisation will recognise — cryptography, identity and access management, patch handling, logging, data flow architecture — albeit applied with constraints they may not have encountered before. The next article picks up the most prescriptive of these: the cryptographic baseline assumed in every European bid, where vendor and operator expectations diverge most quietly.

This article reflects the regulatory and standards landscape at publication. The Cyber Resilience Act's implementing acts continue to be issued through 2026–2027 and may alter specific SBOM and vulnerability handling obligations. Named tools, vendors and platforms are illustrative rather than endorsements. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

The cryptographic baseline assumed in every European bid

15 May 2026 · 11 min read · #compliance #security #industrial #oem-eu-readiness

Late stage of design review. The asset owner's cryptographic reviewer is going through the manufacturer's product security documentation, looking for the algorithm declarations. The document is well-prepared and complete. It lists the protocols the device speaks, the ciphers supported, the key lengths used, the hashing algorithms employed at various points in the firmware.

The reviewer pauses on a particular line. The firmware signature verification — the cryptographic check the controller performs before accepting a firmware update — uses SHA-1 as the hash function and RSA-2048 as the signature algorithm.

The reviewer makes a note. Two issues with that line. First, SHA-1 has been demonstrably broken for collision resistance since 2017 , has been deprecated across every major cryptographic guidance body, and is formally retired by NIST for digital signature use at the end of 2030. The current cryptographic review will not accept it today, regardless of the formal deadline. Second, RSA-2048 sits on the boundary of what is acceptable; NIST SP 800-131A Rev. 2 recommends migration to 3072-bit RSA or to elliptic curve signatures for any device with a service life beyond 2030, which describes essentially every industrial controller being procured in 2026.

The manufacturer's lead engineer is surprised. The SHA-1 was deliberate — chosen for fast computation on the resource-constrained controller eight years ago, never revisited. The RSA-2048 was the standard for the bootloader code at the time, defensible then, on the edge

now. Both choices were reasonable when made. Neither survives a current cryptographic review.

This is the cryptographic baseline conversation, and it is the area in which vendor assumptions and EU expectations diverge most quietly. Most of the architectural arguments in this series surface in early conceptual conversations; cryptography, by contrast, often passes through factory acceptance, gets installed at site, and reveals its mismatches only when a security review is conducted six months into operations. The lender's technical adviser opens a finding. The manufacturer's engineering team faces a firmware update that touches the bootloader, deployable only during planned outages, requiring re-certification under the manufacturer's own quality system. The remediation is expensive and slow. The conversation, had at proposal stage, would have cost almost nothing.

The baseline at the transport level

For network communication, the floor expected in any EU-financed project is TLS 1.3 — the current standard, in widespread deployment since 2018, preferred for all new product designs. TLS 1.2 remains acceptable for legacy migration paths but only with a restricted cipher suite list: forward-secrecy key exchange through ECDHE rather than static RSA, AEAD cipher modes (AES-GCM, AES-CCM, ChaCha20-Poly1305), and removal of all CBC-mode ciphers except where mandatory and properly mitigated against known attacks. SSL of any version — SSLv2, SSLv3 — has been deprecated for over a decade and must not be present, even as a configuration option. TLS 1.0 and TLS 1.1 are similarly out of scope.

For asymmetric cryptography, the floor for new designs is RSA-3072 or elliptic curve cryptography on NIST P-256 (or higher), with Ed25519 increasingly accepted for newer implementations. RSA-2048 is acceptable for legacy products under migration but is on the

boundary of what should be specified for new equipment. For Diffie-Hellman key exchange, ECDHE on standard curves is the expected choice; finite-field DH should be 2048 bits minimum, but EC variants are now standard practice.

For hashing, the floor is SHA-256, with SHA-384 expected for higher security level zones and SHA-512 acceptable. SHA-1 is not acceptable for any cryptographic purpose — neither for signature verification, nor for HMAC, nor for file integrity checking. The formal NIST retirement deadline is end of 2030, but in practice the asset owner's reviewer will flag SHA-1 use today, not in 2030. MD5 has been broken since 2004 and has no legitimate cryptographic use in industrial equipment.

For symmetric encryption, AES is the floor, with AES-128 acceptable for most use cases and AES-256 expected for SL-T 3 zones and above. ChaCha20 is acceptable as an alternative, particularly on resource-constrained devices that struggle with hardware-accelerated AES. 3DES is unacceptable. DES, RC4 and any export-grade cipher are unacceptable and should not be present in firmware even as fallback options.

For digital signatures — the most consequential application of asymmetric cryptography in industrial equipment, because firmware signature verification depends on it — ECDSA on NIST P-256 (or higher) is the standard choice, with Ed25519 increasingly preferred for newer designs and RSA-3072 acceptable. The signature algorithm used by the bootloader to verify firmware updates is one of the cryptographic checks the asset owner's reviewer will examine most carefully, because a compromised firmware signature mechanism compromises every other security control the device has.

These are not parameter choices to be left to deployment configuration alone. The device must support these algorithms in firmware. Weaker algorithms should not be available as fallback options the de-

ployment configuration must remember to disable. Where weak algorithms are present for historical compatibility — a TLS 1.2 server that can still negotiate SHA-1 with old clients, an HMAC implementation that retains MD5 for legacy interoperability — their removal must be planned as part of the product roadmap, not left as a future maintenance topic.

Certificates and PKI integration

The cryptographic algorithms are one half of the baseline. The other half is how the device handles the cryptographic identity infrastructure that the asset owner operates.

A persistent issue with industrial equipment is the self-signed or vendor-rooted certificate baked into the firmware at manufacture. A device that ships with a self-signed certificate and offers no facility for replacement is a device whose cryptographic identity cannot be integrated into the asset owner's public key infrastructure. Every TLS handshake the device performs falls outside the asset owner's chain of trust; every certificate validation requires explicit override; every audit of the cryptographic boundary returns the same finding.

The expectation is that devices accept certificates issued by the asset owner's PKI, deployed through the asset owner's certificate enrollment process. The protocols for that enrollment are standardised — SCEP (Simple Certificate Enrollment Protocol), EST (Enrollment over Secure Transport, RFC 7030) , and increasingly ACME (Automated Certificate Management Environment, RFC 8555) — and the device should support at least one of them. A device that requires manual certificate import through a vendor-specific tool may technically pass acceptance, but it creates operational friction that the asset owner's identity team will surface during onboarding and ongoing operations.

Certificate rotation must be supported. Long-lived certificates — five years, ten years, the device's service life — are increasingly unaccept-

able. The expectation in modern PKI is short-lived certificates, often 90 days or less for connection-level credentials, rotated automatically through the enrollment protocol. A device that supports only certificates with multi-year lifetimes is a device incompatible with the asset owner's PKI policy.

The device must also support certificate revocation in a way the asset owner can verify — either through OCSP (Online Certificate Status Protocol) or through Certificate Revocation List distribution. A revoked certificate that the device continues to trust is a security failure the asset owner cannot tolerate. The verification path must be configurable: pointing at the asset owner's revocation infrastructure, not at the manufacturer's.

Mutual TLS — where both sides of a connection present and verify certificates — is expected for machine-to-machine connections in higher security level zones. Server-side-only TLS is acceptable for some use cases but inadequate for control-system communication where both endpoints need cryptographic identity.

The reference standard pulling this together is IEC 62443-4-2 component requirement 1.8 (public key infrastructure certificates) and CR 1.9 (strength of public key authentication), with the latter's higher security level enhancement requiring hardware-based protection of authentication keys.

Hardware root of trust and secure boot

The cryptographic baseline is increasingly anchored at the hardware level, not just in firmware configuration.

A hardware root of trust — a chip or chip area dedicated to cryptographic operations and key storage, separated from the general application processor — is the foundation. Common implementations include TPM 2.0 (Trusted Platform Module, ISO/IEC 11889), ARM TrustZone with a secure element, dedicated secure cryptographic

chips (Microchip ATECC608, Infineon OPTIGA, NXP A71CH and their successors), and FPGA-based hardware security modules. The role is consistent: provide a tamper-resistant location for storing private keys, performing cryptographic operations, and anchoring the chain of trust for secure boot.

Secure boot is the discipline that uses the hardware root of trust to verify, cryptographically, that every stage of the boot process — bootloader, kernel, root filesystem, application — has been signed by an authorised party before it is allowed to execute. A modified or unsigned image is rejected at the first verification step. The chain of trust extends from the immutable hardware root through every loaded component, with each stage verifying the next before passing control. Tampering at any point breaks the chain and prevents the device from booting.

For SL-T 3 zones and above, hardware root of trust and secure boot are increasingly assumed. A device without these features can be deployed into lower security level zones with compensating controls, but its presence in an SL-T 3 zone requires explicit risk acceptance an asset owner is increasingly unwilling to grant. The trend across the industry is toward hardware-anchored security as standard practice, with the cost of inclusion having fallen to the point where it is no longer a meaningful BOM consideration on most controller-class hardware.

The cryptographic keys held in the hardware root of trust are the device's identity. They are provisioned at manufacture, ideally in a controlled environment with auditable processes. The device's certificate, issued by the asset owner's PKI, is bound to a key pair whose private key never leaves the secure element. This binding is what makes the device's cryptographic identity trustworthy — the private key cannot be extracted, cloned or substituted, even by an attacker with physical access to the device.

Post-quantum and the long view

A topic the cryptographic baseline conversation now touches, where it would not have five years ago, is post-quantum cryptography.

The risk is well-defined. Sufficiently large quantum computers — when they exist, on a timeline still debated but generally placed somewhere between 2030 and 2040 — will break RSA and elliptic curve cryptography. Encrypted data captured today, stored, and decrypted later will be readable. Authenticated connections established today, with logs preserved, will be retroactively impersonable. The "harvest now, decrypt later" attack model is taken seriously by intelligence services and increasingly by lender risk teams underwriting twenty-five-year assets.

NIST completed the first round of post-quantum cryptography standardisation in August 2024, publishing FIPS 203 (ML-KEM, formerly CRYSTALS-Kyber, for key encapsulation), FIPS 204 (ML-DSA, formerly CRYSTALS-Dilithium, for digital signatures), and FIPS 205 (SLH-DSA, formerly SPHINCS+, for hash-based signatures). The European cryptographic community has broadly endorsed these algorithms. The transition is now an engineering matter rather than a standards matter.

The expectation for new industrial equipment in 2026 is not that products implement post-quantum cryptography today — that would be ahead of mainstream practice and may not interoperate with the rest of the ecosystem. The expectation is that manufacturers have a credible roadmap for it, that they are tracking the standardisation outcomes, and that their cryptographic agility — the ability to swap algorithms without redesigning the product — supports a future migration. Hybrid TLS 1.3 implementations, combining a classical algorithm and a post-quantum algorithm in a single handshake, are increasingly seen in serious deployments and are likely to be the migration path for industrial systems.

A device whose cryptographic implementation is hard-coded — algorithms compiled into firmware with no facility for replacement, key sizes fixed, ciphers immutable — is a device that cannot make the post-quantum transition without firmware replacement. For equipment with a twenty-five-year service life, this is a material commercial risk the lender's risk team will price into their assessment.

At proposal stage

A manufacturer's bid that includes a cryptographic configuration document — listing the algorithms supported, the protocols implemented, the key management approach, the hardware root of trust if present, the secure boot chain if present, the certificate enrollment protocols supported, and the post-quantum roadmap — is a bid that has anticipated the conversation. The lender's technical adviser reviews the document, identifies any specific concerns, and the conversation proceeds.

A manufacturer's bid that does not address cryptography, or that addresses it generically with phrases like "industry-standard encryption" or "secure protocols supported", signals that the team has not yet considered the specifics. The cryptographic review will surface gaps; the gaps will require firmware modifications; the firmware modifications will need to be planned, tested and deployed before commissioning. The cost of the modifications is rarely large in absolute terms. The schedule impact, if surfaced late, can be considerable.

The cryptographic baseline is, in some ways, the most universal of the topics in this series. It does not depend on regulation in the way the disclosure programme or the bill of materials do. It does not depend on the project structure in the way the network and substation boundaries do. It is a property of the equipment, applicable to every deployment, every customer, every jurisdiction. A manufacturer that

builds cryptographic discipline into their product line is a manufacturer that has reduced friction across their entire market — EU-financed and otherwise.

The next article picks up the engineering function that cryptography supports but does not, on its own, deliver: the identity and access model for engineers, service accounts and machine-to-machine connections, where the manufacturer's habit of named-team access through shared credentials meets the asset owner's expectation of named-person access through their identity infrastructure.

This article reflects the cryptographic landscape at publication. NIST guidance on algorithm deprecation continues to evolve, particularly around the SHA-1 retirement deadline and post-quantum migration timelines. References to commercial secure element vendors are illustrative rather than endorsements. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

Bring your engineers, not your accounts

15 May 2026 · 11 min read · #compliance #security #industrial #oem-eu-readiness

An audit trail review, three months after commissioning. The asset owner's security operations team is reconstructing a series of configuration changes made to a turbine controller during the previous week. The changes were authorised — they appear in the change management system, approved, with clear engineering justification. The session was recorded by the secure remote access broker; the screen capture shows what was done.

What the team cannot reconstruct is who did it.

The session was logged into using an account called `svc_oem_eng`. The account is shared across the manufacturer's service organisation — perhaps eight named engineers have access to its password, distributed via the manufacturer's own internal credential management. The session recording shows the screens, but does not show which of the eight engineers was at the keyboard. The change management ticket lists "manufacturer service team" as the actor. The audit trail, in the strict sense the asset owner's compliance function needs, does not exist.

This is the identity model the manufacturer brought to the project. Perfectly normal in their own internal operations, perfectly compatible with the way their service organisation has always worked, perfectly insufficient for an EU-financed project under NIS2 .

The asset owner's expectation is straightforward. Every action against an OT asset must be attributable to a named individual person, authenticated through the asset owner's identity infrastructure, authorised through the asset owner's privileged access management

process, audited under the asset owner's logging discipline. The manufacturer's engineers do not bring their own identities into the project; they receive identities from the asset owner. Service accounts for application-to-application communication exist, but they are managed and rotated by the asset owner's secret store, not held as static credentials in the manufacturer's tools. Hard-coded passwords in firmware are not configuration choices; they are delivery defects.

The principle is named-person access, time-bound, audited per person, with separation between the human identity surface and the machine-to-machine identity surface. The shorthand the security community sometimes uses is zero trust; the underlying engineering predates the label by a decade and is now mainstream in serious asset owner organisations.

What the asset owner provides

The identity infrastructure the manufacturer's engineers will integrate into is a defined stack.

The asset owner runs an enterprise identity and access management platform — typically built around Microsoft Entra ID (the platform formerly known as Azure Active Directory), Okta, Ping Identity, or a similar enterprise provider — and a privileged access management platform — typically CyberArk, BeyondTrust, Delinea, or HashiCorp Vault depending on the asset owner's preference. These platforms manage user identities, group memberships, authentication factors, session brokering for privileged access, and the audit logs for all authentication and authorisation events.

The manufacturer's engineers receive identities in this infrastructure. Each engineer has a named account: not `svc_oem_eng` but a personal identifier — Wei Chen, Maria Andersson, Hiroshi Tanaka. The account is provisioned through a documented onboarding process that in-

cludes background check completion, training on the asset owner's procedures, and acknowledgement of the asset owner's acceptable use policy. Multi-factor authentication is configured through the asset owner's MFA infrastructure — phone app, hardware token, FIDO2 key — not through the manufacturer's.

The account has a defined validity period. It does not exist indefinitely. The default expiry is the end of the engineer's certification on the equipment, or the end of their employment with the manufacturer, or a specific calendar date — whichever comes soonest. Renewal is a deliberate act, requiring confirmation that the engineer is still certified and still requires access.

For privileged access — the access that allows actual changes to the controllers and the SCADA, rather than read-only observation — the engineer authenticates through the privileged access management platform, which brokers session access to specific targets for specific durations under specific authorities. Standing privileged accounts on the target systems do not exist; the engineer's privileged access is established when needed and revoked when finished, with the brokered session model described earlier in this series .

IEC 62443-3-3 organises the requirements for identification and authentication under foundational requirement 1 (identification and authentication control), with the component-level requirements appearing in 62443-4-2 under the same numbering. The asset owner's infrastructure is designed to meet these requirements at SL-T 3 or above; the manufacturer's equipment is expected to be capable of integrating with that infrastructure rather than substituting for it.

Why federation does not solve this

A common manufacturer's proposal at this point in the discussion is federation — that the engineers continue to authenticate against the manufacturer's own Active Directory or identity provider, with trust

established between the manufacturer's IDP and the asset owner's IAM through SAML or OpenID Connect. This is operationally simpler for the manufacturer; their engineers continue to use familiar credentials, their MFA continues to work as deployed, their identity lifecycle continues to flow through the manufacturer's HR system.

For an EU-financed project, federation is generally not accepted at the OT boundary. Three reasons.

First, the asset owner cannot validate the manufacturer's identity practices. They do not know how the manufacturer authenticates a new joiner, how they handle a leaver, how they protect against credential theft, how their MFA is configured, how compromise is detected and responded to. Trusting a federated identity is trusting the entire identity infrastructure of the federating party. The asset owner has no contractual standing to audit that infrastructure deeply, no visibility into its day-to-day operation, no ability to detect compromise in time to respond.

Second, the audit trail breaks at the federation boundary. The asset owner can record that "an identity from the manufacturer's IDP, claiming to be Wei Chen, authenticated and accessed the controller". They cannot independently verify that the identity actually was Wei Chen, rather than someone who had compromised Wei Chen's credentials at the manufacturer. Under NIS2 incident reporting obligations, the asset owner needs an audit trail they can stand behind without dependency on a third party's identity practices.

Third, the leaver process becomes asynchronous. When Wei Chen leaves the manufacturer, the manufacturer's HR system processes the departure, the IDP eventually removes the account, the federation eventually reflects the removal in the asset owner's IAM. The window between Wei Chen's last day and access removal at the asset owner's end is operationally significant for any privileged role. Federation amplifies this window; direct provisioning collapses it.

The accepted pattern is direct provisioning: the engineer has a named account in the asset owner's IAM, provisioned through the asset owner's onboarding, deprovisioned through the asset owner's leaver process, authenticated through the asset owner's MFA, audited through the asset owner's logs. The manufacturer's own identity infrastructure is irrelevant to the project's access architecture.

Federation may still be acceptable for non-OT systems — corporate collaboration tools, ticketing platforms, document repositories — where the security boundary is less sensitive and the audit obligations are different. The OT boundary specifically is where federation reaches its limit.

Service accounts and machine-to-machine identity

The pattern that applies to human engineers also applies, in a different form, to service accounts — the identities used by applications, scripts, and machine-to-machine integrations.

A turbine manufacturer's condition monitoring system pulls telemetry through a specific account on the historian. A patch deployment script logs into the controllers using a service account. A diagnostic tool authenticates to the SCADA's API. Each of these is a non-human identity, with credentials that must be managed somehow.

The historical practice — credentials hard-coded in configuration files, scripts, or firmware — is not acceptable. Hard-coded credentials cannot be rotated without re-deploying the calling system, cannot be revoked without breaking the integration, cannot be audited per use, and frequently end up in source code repositories or backup archives that nobody intended to protect. A device that ships with hard-coded credentials in its firmware is delivering a credential into the asset owner's environment without the asset owner's control.

The expectation is that service accounts are managed by the asset owner's secret store — HashiCorp Vault, CyberArk Conjur, or similar

— with rotation managed by the store, retrieval by authorised consumers at runtime, and audit of every retrieval. Applications retrieve their credentials at startup or on a defined cadence, rotate them on schedule, and never persist them outside the store's protected paths.

For machine-to-machine connections where higher assurance is required, certificate-based authentication replaces password-based service accounts entirely. Each application or device has a certificate issued by the asset owner's PKI (the topic of the previous article in this series), with the private key held in a hardware security module or equivalent. Authentication is mutual TLS; the certificate's identity is the authoritative one. Rotation of certificates replaces rotation of passwords. Audit of certificate issuance and use is part of the PKI infrastructure rather than a separate logging exercise.

Modern machine-identity frameworks — SPIFFE and SPIRE , increasingly seen in serious deployments — formalise this approach by issuing short-lived cryptographic identities to workloads and services through a workload identity attestation process. The frameworks are not yet universally deployed in OT, but the direction is clear: certificate-based machine identities, short-lived, automatically rotated, audited per use.

A device that supports only password authentication for its service interfaces is a device with an architecture limitation that will be raised at design review. Devices that support certificate-based authentication, with the certificate enrollment protocols from the cryptographic baseline article, integrate cleanly into modern identity infrastructure.

What this means for the manufacturer's organisation

The shift from team-pool access to direct named-person access has organisational consequences the manufacturer should anticipate.

The named engineers must be disclosed in advance. The manufacturer cannot operate a model in which "whichever engineer is available"

handles a service call; specific engineers must be identified, vetted, onboarded, and trained on the asset owner's procedures before they can access the project's systems. For a large service organisation, this typically means a designated pool of certified engineers for the project — perhaps a dozen people, perhaps fewer — rather than an open roster.

The pool must be maintained. When an engineer leaves the manufacturer, the asset owner needs to be notified promptly. When a new engineer is added to the pool, the onboarding process runs again. When an engineer's certification on a piece of equipment expires, their access is suspended until recertification. The manufacturer's HR and certification systems need to interface with the asset owner's identity lifecycle — either through formal integration or through reliable manual notification — and the responsibility for keeping the pool current rests with the manufacturer, not the asset owner.

The manufacturer's internal practice of using shared service accounts for internal tooling does not transfer to the asset owner's environment. Inside the manufacturer's own systems, shared accounts may be efficient and operationally acceptable. Inside the asset owner's OT environment, they are not. The manufacturer's engineers will in practice have two identity surfaces: their internal manufacturer identity for their own tools and systems, and their project-specific named identity for any access to the deployed assets. Maintaining the separation between these two surfaces is part of the service organisation's discipline rather than a one-off setup task.

This duality is not unique to EU-financed projects; it is the direction industrial identity management has been moving across the global industry for the last decade. EU regulation has accelerated and codified the practice, but the underlying engineering — named-person access, time-bound credentials, certificate-based machine-to-machine authentication — is mainstream rather than exotic.

At proposal stage

A manufacturer's bid that addresses identity and access — that proposes a named pool of certified engineers for the project, describes how their credentials will be onboarded into the asset owner's IAM, commits to certificate-based authentication for the equipment's service interfaces, and demonstrates that no hard-coded credentials are present in shipped firmware — is a bid that has anticipated the conversation. The conversation that follows is operational: how the onboarding workflow integrates with the manufacturer's service rotation, what the recertification cadence will be, how leavers will be communicated.

A manufacturer's bid that proposes "secure VPN access for our service organisation using our corporate credentials" signals that the team is operating from the previous decade's identity model. The model is not unbridgeable — the same engineers can be onboarded into the asset owner's IAM, the same equipment can be reconfigured for certificate-based authentication, the same service accounts can be migrated to a managed secret store — but the work to close the gap needs to start before contract signature.

The deeper observation, similar to the one made in the cryptographic baseline article, is that this discipline pays off across the manufacturer's entire market. Named-person access, time-bound credentials, certificate-based authentication, and managed service accounts are not EU-specific requirements; they are the direction every serious asset owner's identity infrastructure is moving. A manufacturer that has adapted to deliver into one EU-financed project has adapted to deliver into the next one, and into the high-end of the global market more broadly.

The next article picks up the topic identity supports but does not by itself address: the patch delivery contract between the manufacturer and the asset owner, where the historical practice of vendor-pushed

automatic updates meets the operator's expectation of controlled, scheduled, signed deployments arriving through a documented release process.

This article reflects the identity and access management landscape at publication. References to commercial IAM, PAM and secret store platforms are illustrative rather than endorsements; the underlying engineering principles apply across vendor choices. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

Patches arrive on the owner's schedule, not yours

15 May 2026 · 11 min read · #compliance #security #industrial #oem-eu-readiness

Monday morning operations review. The asset owner's commissioning team is going through the previous week's alarms when one of the engineers notices something unusual. Three of the wind turbines logged a controlled restart at 03:42 on Saturday. The restart was not in the change management system. The control room saw the restarts, recorded them as routine, watched the turbines come back online cleanly.

Investigation begins. The audit logs on the controllers show a firmware update was applied at 03:38, four minutes before the restart. The update package signature is valid. The update came from a connection the asset owner's security team is now examining: an outbound TLS session from the controller to a hostname that resolves to the manufacturer's update infrastructure. The connection was permitted by the firewall rule that allowed the controller to report telemetry — the rule that, on close reading of the manufacturer's documentation, was supposed to be outbound-only for diagnostic data, not bidirectional for software delivery.

The asset owner's security team escalates. The firmware update is benign in itself — a routine quarterly maintenance release, well-tested, no security implications either positive or negative. The mechanism is not benign. The manufacturer's controllers have just demonstrated that they can receive and apply software changes from outside the asset owner's control. Three out of fifty turbines applied the update; the other forty-seven did not, for reasons the manufacturer cannot immediately explain. The asset owner now has a fleet in mixed configurations, an audit trail with software changes that did not pass

through change management, and a regulatory disclosure to consider under NIS2 .

The principle the asset owner expects is brief, and the manufacturer's bid did not address it. Software changes to the deployed equipment happen on the asset owner's schedule, under the asset owner's change management, with the asset owner's approval. The manufacturer's role is to make patches available — signed, documented, tested, accompanied by the information required for the asset owner to make a deployment decision. The deployment itself is the asset owner's act, not the manufacturer's.

Why OT patching differs from IT patching

A reasonable manufacturer's reaction at this point in the conversation is that the auto-update mechanism is industry-standard, supported by every major cloud platform, used by everything from consumer phones to enterprise software. It is. The difference is the operating context.

In information technology, an update can be hot-deployed during business hours, the affected user retries their action, the inconvenience is brief, the rollback is a software re-image. In operational technology, an update typically requires a restart, the restart requires coordination with the grid operator (because the plant briefly stops producing), the affected behaviour might have safety implications, the rollback might require physical access to the controller. A bad update applied to fifty turbines simultaneously can take a plant offline for the rest of the day, with revenue impact, grid stability consequences and contractual penalties under the connection agreement.

The asset owner is also the regulated entity. Under NIS2, the operator is responsible for the cybersecurity posture of the asset and for incident notification if something goes wrong. An update applied without the operator's knowledge that subsequently causes an inci-

dent is an incident the operator must report and explain to the national CSIRT, while not having had control over the change that caused it. This is a regulatory position no operator is willing to occupy.

The Cyber Resilience Act , in Article 13 (via Annex I Part II point 8), requires manufacturers to provide security updates promptly and free of charge during the declared support period. The Act does not require those updates to be pushed automatically. The deployment mechanism is the operator's to choose. The mature operator chooses controlled deployment because the consequences of uncontrolled deployment are exactly what the auto-update mechanism cannot mitigate.

IEC 62443-4-1 organises the manufacturer's responsibilities for security update management under the practice set known as SUM, covering qualification (SUM-1), documentation (SUM-2), dependent component documentation (SUM-3), delivery (SUM-4), and timely delivery of security patches (SUM-5). The standard makes the manufacturer responsible for making patches available with the necessary documentation and assurance evidence. It does not make the manufacturer responsible for deploying them — that responsibility, in a properly architected service relationship, sits with the asset owner.

The patch delivery contract

What replaces auto-update is a structured contract between the manufacturer and the asset owner. The contract has eight stages.

Release. The manufacturer publishes the patch in their controlled release infrastructure — a download portal that requires authentication, accessible to authorised asset owner representatives. Public download URLs, repositories accessible without credentials, and update servers that controllers reach autonomously are not part of this model. The manufacturer's release infrastructure is the point at

which a new version becomes available; it is not the point at which it is deployed.

Notification. The asset owner is notified through the channel agreed in the contract. For security patches, this is typically the manufacturer's PSIRT advisory channel , with the advisory cross-referenced to the patch release. For feature and maintenance updates, the notification channel is usually the manufacturer's customer success or technical account management function. The notification specifies what the patch contains, why it is being released, what the assessed criticality is, and what the recommended deployment window looks like.

Delivery. The patch is delivered as a signed artefact — a firmware image, a configuration package, an application installer — with a cryptographic signature using the manufacturer's release signing key. The asset owner verifies the signature against a public key established at contract signature and held in their PKI. A patch whose signature does not verify is not deployable, regardless of how urgent the manufacturer's notification claims it to be.

Documentation. Each release includes a documentation set: release notes describing what changed and why; known issues and their workarounds; regression test evidence describing what was tested by the manufacturer and how; rollback procedure describing how to revert if the deployment causes issues; and a software bill of materials diff showing which components changed at what versions. The documentation set is what enables the asset owner's change advisory board to make an informed deployment decision rather than relying on the manufacturer's assurance alone.

The CSAF format — Common Security Advisory Framework, an OASIS standard — has emerged as the machine-readable structure for security advisories and patch metadata. Manufacturers that publish CSAF-format advisories integrate cleanly into the asset owner's vulnerability management tooling. Those that publish PDF advisories re-

quire the asset owner's team to transcribe the information, which slows the process and introduces error.

Testing. For SL-T 3 zones and above, the asset owner typically maintains a staging environment that mirrors the production configuration for representative equipment. Patches are deployed to staging first, run for an agreed observation period, evaluated for regression and unexpected behaviour, and only then approved for production. For lower security level zones, the asset owner may accept the manufacturer's regression evidence as sufficient and skip the staging step. The decision is the asset owner's to make.

Approval. The asset owner's change advisory board reviews the patch package, the documentation, the staging results if applicable, and the proposed deployment plan. Approval is granted under the asset owner's change management process, with the resulting change ticket carrying the audit trail.

Deployment. The patch is applied during a scheduled outage window, coordinated with operations and the grid operator. The deployment mechanism is the asset owner's, using their authenticated session against the controller through the secure remote access broker. The manufacturer's engineering team may be in attendance, but they are not the actors; their role is to advise if anything unexpected occurs.

Post-deployment verification. The asset owner verifies the patch has been applied successfully, runs functional tests against the updated equipment, monitors operational behaviour for the period agreed in the change ticket, and confirms successful completion in the change management system. Failed deployments are rolled back using the manufacturer's documented procedure.

This is the contract. It is more elaborate than auto-update because the consequences of getting it wrong are larger. It is also the discipline that every mature asset owner's change management operates

under, regardless of EU regulation, because controlled change is the foundation of operational reliability.

Emergency patches and silent updates

A separate but related topic is what happens when a critical vulnerability is announced — exploited in the wild, scored at the top of the severity scale, demanding immediate response.

Emergency patches follow an accelerated path through the same contract, not a parallel path outside it. The PSIRT advisory is issued; the asset owner's PSIRT-watching function picks it up; the asset owner's emergency change process activates; a same-day or next-day deployment is scheduled if the risk warrants it; the manufacturer's release infrastructure provides the signed artefact and the documentation; the asset owner deploys under their authority. The compressed timeline does not change who performs the deployment or under whose authorisation.

The manufacturer does not deploy emergency patches on their own authority. Even when the manufacturer has the technical capability to do so — because the controllers can receive updates through some mechanism, perhaps the same one the auto-update incident in the opening scene revealed — the act of deploying without asset owner approval is a contractual violation in EU-financed projects. The asset owner's risk assessment may decide that the vulnerability is genuinely emergency-grade and approve deployment within hours; that decision is theirs.

The "no silent updates" principle has technical implications. Controllers must not accept software updates from external sources without explicit local authorisation. Telemetry channels must be strictly outbound, with the firewall rules and unidirectional gateway architecture enforcing the constraint at the network level. Update mechanisms must require authentication using the asset owner's creden-

tials or PKI, not credentials the manufacturer holds independently. Local update applications — engineer at the cabinet with a laptop, USB drive with a firmware image — must follow the change management process, with the deployment tracked the same way a remote deployment would be. Software composition must be verifiable at any time, typically through the SBOM and a cryptographic attestation of the running firmware against the signed release.

A device that supports auto-update but allows it to be disabled is not adequate. A device that supports auto-update enabled by default, configurable to off at deployment, is also not adequate — the default configuration is what the asset owner's commissioning team has to remember to change, and the auto-update mechanism is what their security team has to remember to monitor for re-enablement. The expectation is that the equipment is delivered with auto-update structurally disabled — either not implemented in the project variant of the firmware, or implemented only against a release server that the asset owner controls.

At proposal stage

A manufacturer's bid that addresses the patch delivery contract — that describes the release infrastructure, the signature mechanism, the notification channel, the documentation set including SBOM diff, the support for rollback, the absence of auto-update mechanisms in the project variant — is a bid that has anticipated the conversation. The conversation that follows is operational: the cadence of expected releases, the channel for emergency advisories, the staging environment specifications, the rollback procedures specific to this equipment.

A manufacturer's bid that proposes "automatic firmware updates pushed from our cloud for security and feature improvements" signals an architectural model that the asset owner cannot accept. The

work to close the gap involves disabling the auto-update path, providing a signed-artefact delivery mechanism, building or extending the manufacturer's release infrastructure to support authenticated download, and integrating with the CSAF format for advisories. The work is the work of weeks to months depending on the manufacturer's current state, not a procurement footnote.

The deeper observation is that the patch delivery contract is the topic where industrial operational technology has lagged behind general information technology in some ways and led in others. Lagged because the deployment discipline is less automated than IT, where patches can be hot-deployed and rolled back without coordination with the grid operator. Led because the change management discipline is more rigorous than IT, with engineering review, regression testing, and post-deployment monitoring that IT often skips. The patch contract described here combines the best of both — rigorous change management with modern release infrastructure, signed artefacts, machine-readable advisories, and structured rollback. A manufacturer who delivers patches this way has built capability that pays off across their global service organisation, not only in EU-financed projects.

The next article picks up the operational telemetry that flows in the opposite direction from patches: the logging and monitoring discipline where the asset owner's security operations centre, not the manufacturer's, is the system of record for security events.

This article reflects the regulatory and standards landscape at publication. The Cyber Resilience Act's implementing acts continue to evolve through 2026–2027 and may alter specific patch delivery and security update obligations. References to IEC 62443-4-1, CSAF and

related standards may be superseded by revisions. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

Logs ship in formats someone else's SOC can read

15 May 2026 · 12 min read · #compliance #security #industrial #oem-eu-readiness

Three-fourteen in the morning. The asset owner's security operations centre receives an alert: anomalous outbound connection attempts from a SCADA workstation in the wind farm control room. The analyst on duty pulls up the device logs to investigate. The SCADA application's log shows the timestamp that the alert correlated to, then a single line: `NetworkException: connection failed`. No source process, no destination, no protocol, no port, no user, no return code. Just the message and the time.

The analyst escalates to a senior analyst, who pulls the firewall logs for context. The firewall shows three connection attempts, two seconds apart, all blocked — outbound to a hostname that resolves to an IP address in a jurisdiction the asset owner's policy flags as restricted. The source was the SCADA workstation. The initiating process is identified at the network layer but not the application layer; the firewall sees the parent of the OS network stack, not the executable that opened the socket.

The senior analyst calls the manufacturer's emergency line. The on-call engineer takes the report. Twenty minutes later they come back. Yes, they have seen this. The manufacturer operates a fleet-wide telemetry system that monitors network behaviour on their SCADA installations across multiple customers. The same anomalous pattern appeared at three other sites the previous month, traced to a misconfigured update check in a third-party library bundled with the SCADA application, patched in the latest release. They had not communicated the issue because their support contract did not specifically require them to.

The asset owner's SOC has spent ninety minutes investigating an event the manufacturer already knew about. The manufacturer's monitoring picked it up, diagnosed it, fixed it in the product, and did not tell the operator whose plant they were watching. The architectural failure is not in either side's monitoring. It is in the assumption that two parallel monitoring systems was an acceptable architecture.

The principle the asset owner expects is direct. The security operations centre is the system of record for security events affecting the deployed equipment. Not the manufacturer's monitoring centre, not the manufacturer's fleet analytics, not the manufacturer's support ticket system. Every security-relevant event the equipment generates must reach the asset owner's SIEM (Security Information and Event Management platform) in a format the SIEM can consume, with the fields necessary for investigation, time-synchronised to the asset owner's authoritative time source, with an audit trail that is tamper-evident and that the manufacturer cannot unilaterally delete or rotate.

The manufacturer's own monitoring may continue to exist — there are good reasons for it to exist, including fleet benchmarking and predictive maintenance — but it is not a substitute for the operator's monitoring. The two systems serve different purposes, sit under different governance, and are accountable to different parties.

What the SIEM needs

The asset owner's SIEM is the integration point for security event data across the entire OT environment. It correlates events from firewalls, intrusion detection systems, jump hosts, secure remote access brokers, identity platforms, network switches, and the OT devices themselves. To do that work, it needs events in a specific shape.

Standard format. The most widely supported formats are Syslog (RFC 5424 , with structured data extensions), the Common Event Format originally from ArcSight, and the Log Event Extended Format origi-

nally from IBM QRadar. OpenTelemetry is increasingly accepted as a modern alternative, particularly for newer deployments. The choice between them is the asset owner's, but the device must support at least one. Proprietary log formats that require custom parsers, manufacturer-specific dashboards as the only viewing mechanism, or binary log files that require manufacturer-supplied tooling to interpret — none of these integrate cleanly. The transport itself should be secure: Syslog over TLS (RFC 5425) is the floor, not plain Syslog over UDP that was acceptable two decades ago.

Required fields. For each security-relevant event, the SIEM needs to know when it happened (timestamp at millisecond resolution, in UTC, sourced from the agreed time reference), what happened (a structured event type from a documented taxonomy), who was involved (the identity that performed the action, the device that generated the event, the source and destination if it is a network event), what the outcome was (success, failure, blocked, allowed), and any context that distinguishes this event from similar events (session identifier, ticket reference, correlation token). Events that lack these fields cannot be correlated effectively with events from other sources, which is the entire reason the SIEM exists.

Documented event taxonomy. The manufacturer should publish, as part of the product documentation, the complete list of security events the device generates — what each event type means, when it fires, what fields it includes, how it differs from similar events, what severity it should be treated as. Without the taxonomy, the SIEM team is reverse-engineering event meanings from observed behaviour. This works for the loudest event types and misses the subtler ones. A well-documented event taxonomy is often the most useful single artefact a manufacturer can provide to a deploying asset owner.

Time synchronisation. Every device emits timestamps. Those timestamps must be synchronised to the asset owner's authoritative time

source — typically a GNSS-disciplined master clock distributing NTP or PTP, as covered earlier in this series for the substation side . A device whose clock drifts by minutes makes correlation difficult; a device whose clock drifts by hours makes investigation impossible. The agreed time source, the protocol, and the acceptable drift tolerance are part of the operational specification, and the manufacturer's equipment must accept the time source the asset owner provides rather than relying on a hard-coded NTP pool or, worse, a vendor-internal time source.

Audit logs as a distinct category. Audit logs — records of who did what against the device — are a specific subset of the device's log output, and they have stricter handling requirements than operational logs. They must be tamper-evident, typically through cryptographic hash chaining or write-once storage on the device. They must be forwarded to the SIEM in real time, not stored locally and pulled on demand. They must not be deletable or rotatable by an unprivileged user, and the manufacturer's service engineers must not be able to delete or rotate them through any service interface, including service-mode access during maintenance. The audit log is the evidence base for incident investigation; its integrity is what makes the investigation defensible to a regulator.

IEC 62443-3-3 organises the system-level requirements for audit logging under foundational requirement 2 (use control), with SR 2.8 covering auditable events, SR 2.9 covering audit storage capacity, SR 2.10 covering response to audit processing failures, and SR 2.11 covering timestamps. The component-level equivalents in 62443-4-2 are addressed to the manufacturer directly.

Why the manufacturer's monitoring is not a substitute

The manufacturer's own monitoring serves their interests. They benchmark fleet performance, identify common fault patterns, im-

prove their products through observation, support their service operations. None of these activities is bad. None is sufficient.

The asset owner's monitoring serves the asset owner's interests. They detect threats specific to their environment, correlate events across systems the manufacturer cannot see, comply with their regulatory reporting obligations, and maintain the audit trail their lender requires. The two activities overlap in some areas — both observe the device's behaviour, both might notice the same anomaly — but they answer different questions, are accountable to different parties, and produce different artefacts.

Three reasons the two systems cannot be substituted.

The asset owner cannot rely on the manufacturer to share information about events that affect the asset owner's equipment. The manufacturer may have commercial reasons to not share. They may not consider an event significant enough to mention. They may be on a different time cycle for analysis and notification. The asset owner's regulatory obligations — incident notification under NIS2 within 24 hours of becoming aware — assume the asset owner becomes aware through their own infrastructure, not through a third party's selective disclosure.

The SIEM correlates events across systems. A login event from the identity platform, a connection event from the firewall, a process event from the controller, a configuration change event from the engineering workstation — together these tell a story that no single system tells alone. The manufacturer's monitoring sees only the device behaviour; it cannot correlate against the other systems the SIEM integrates with, because it does not see them.

The audit trail must be the asset owner's. Under NIS2 and the broader EU framework, the asset owner needs to be able to evidence what happened, when, on what authority, in their own systems, without de-

pendency on a third party. The manufacturer's logs, held in the manufacturer's monitoring centre, accessible only through the manufacturer's tooling, do not meet this standard. The audit trail must be in the asset owner's hands.

The constructive arrangement is that both monitoring systems exist, with clear scope and a defined information-sharing protocol. The manufacturer monitors fleet behaviour for their own purposes, anonymised where appropriate, with the analytics they need for product improvement. The asset owner monitors all security-relevant events from the deployed equipment in their SIEM. The manufacturer is contractually required to share, in real time, any event their monitoring detects that the asset owner's monitoring would have detected if it had equivalent visibility — meaning anything security-relevant, anything anomalous, anything indicative of compromise. The default position is that information flows from the manufacturer to the asset owner promptly, with the manufacturer's monitoring serving as a supplementary sensor rather than a sealed channel.

False positives and instrumentation gaps

A practical observation worth noting. Industrial firmware is often log-noisy in unhelpful ways. The device emits log messages for routine events that, in modern IT systems, would never be logged at the security event level — successful self-tests, normal protocol handshakes, scheduled tasks completing on schedule. These events flood the SIEM, create false positives in the alert rules, and obscure the events that actually matter. The first weeks of any new deployment involve substantial SIEM tuning to suppress the noise without losing signal.

The remedy is on both sides. The device's logging configuration should be tunable — by event severity, by event category, by source module — so the asset owner's SOC team can suppress noise without losing signal. The SIEM's correlation rules should be authored specif-

ically for the device's event taxonomy, not generic patterns assumed from other equipment. Both adjustments take work. They are normal SOC tuning activities, not signs that the integration is broken.

A different category of problem comes from inadequate device instrumentation — events that should fire and do not, or that fire with insufficient information. A login event without the source IP. A configuration change event without the changed parameter. A network event without the destination. A privilege escalation event that does not specify the privilege being escalated to. Each is a gap in the SIEM's investigative capacity. The remedy here is product improvement: the asset owner raises the gap with the manufacturer, the manufacturer addresses it in a firmware release through the patch contract described in the previous article . Over time the device's logging maturity improves, the false positive rate falls, and the SOC integration becomes routine. The first six months of operation are typically the noisiest; the second six months see substantial improvement.

The asset owner's SOC team is, in this sense, a continuous reviewer of the manufacturer's product quality. They see, in their alert rates and their investigation hours, where the device is well-instrumented and where it is not. Manufacturers who treat the asset owner's SOC findings as product feedback — incorporating instrumentation improvements into the development backlog — see their equipment become easier to integrate over time and their SOC integration cost fall accordingly. Manufacturers who treat the findings as deployment-specific configuration issues to be worked around locally never improve at the source.

At proposal stage

A manufacturer's bid that addresses logging — that specifies the supported formats (Syslog with structured data, CEF, LEEF, or OpenTelemetry), commits to a documented event taxonomy delivered with

the equipment, supports configurable log severity and category filtering, supports time synchronisation to the asset owner's authoritative source, and confirms that audit logs are tamper-evident and forwardable in real time — is a bid that has anticipated the conversation. The conversation that follows is operational: SIEM integration testing during factory acceptance, event taxonomy review with the asset owner's SOC team, log volume estimation for SIEM sizing, alert rule authoring against the documented taxonomy.

A manufacturer's bid that proposes "comprehensive logging available through our cloud monitoring dashboard, accessible to the customer through a customer portal" signals the previous decade's model. The model is incompatible with the asset owner's SIEM-as-system-of-record requirement, and the work to close the gap involves either reconfiguring the device to emit standard-format logs in real time over a secure transport, or building a secondary export path that achieves the same outcome — typically the latter, because the device firmware changes are slower.

The deeper observation is that monitoring discipline pays off across the manufacturer's entire market. Standard log formats, documented event taxonomies, tamper-evident audit logs, time synchronisation, real-time forwarding — these are not EU-specific requirements; they are the direction every serious asset owner's security operations infrastructure is moving. A manufacturer who delivers monitoring this way has reduced friction across their global service organisation and has positioned themselves to integrate with the customer's tooling rather than requiring the customer to integrate with theirs. The shift from "log into our portal to see what your equipment is doing" to "your SIEM is the system of record, we contribute structured events to it" is, in the long run, the easier model for the manufacturer too — fewer customer portals to maintain, fewer parallel monitoring stacks to support, clearer information-sharing protocols when things go wrong.

The next article picks up the topic that monitoring exposes more clearly than any other: where the data the manufacturer collects actually goes, who has access to it, under whose jurisdiction it sits, and whether the architecture as built can survive the lender's transfer impact assessment under the GDPR's cross-border data provisions.

This article reflects the regulatory and standards landscape at publication. References to RFC 5424, RFC 5425, IEC 62443-3-3, and commercial logging formats (CEF, LEEF) are stable but may be supplemented by newer standards such as OpenTelemetry as deployment patterns evolve. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

The data lands somewhere. The lender wants to know where

15 May 2026 · 11 min read · #compliance #security #industrial #oem-eu-readiness

The data protection officer at the EU-based lender's project finance team is reviewing the manufacturer's proposed architecture. She asks one question, in writing, to the manufacturer's bid team:

"When the controller transmits its quarterly performance summary to the manufacturer's analytics platform, what is the physical destination of that data — the data centre, the city, the country, the legal entity that operates the data centre, the legal entity that holds the encryption keys, and the legal entity that has root access to the database?"

The manufacturer's response, several days later, is reassuring in tone. The data is encrypted in transit and at rest. It is held in a secure cloud environment operated by a tier-one cloud provider. Access is restricted to authorised personnel. The manufacturer adheres to international data protection standards.

The data protection officer writes back. "I asked six specific questions. None of them was answered. Please provide the six pieces of information requested."

The second response answers four of the six. The data is hosted in the manufacturer's home jurisdiction. The data centre is operated by a national cloud provider with whom the manufacturer has a long-standing relationship. The legal entity that holds the data is the manufacturer's local analytics subsidiary. The encryption keys are managed by the cloud provider through their key management service. The two remaining questions — the city and the legal entity with root access — are deferred to a later response.

The data protection officer reads the response twice. She has, in those four answers, the basis for declining to recommend the bid. The data flow as described does not survive a Schrems II analysis. The cloud provider, operating in the manufacturer's home jurisdiction, is subject to national surveillance and intelligence law that the European Union has not assessed as adequate. The encryption keys are held by an entity within the same jurisdiction, which means they are accessible to that jurisdiction's authorities under domestic law. The cloud provider's home government has lawful access to the data, even encrypted, under conditions that exceed what the GDPR considers necessary and proportionate.

This is the cross-border data flow problem.

The legal framework was set out in the regulatory reference piece earlier in this series . The articles relevant to the conversation are GDPR Articles 44 to 49 (transfers to third countries), the Schrems II judgment of July 2020 (the transfer impact assessment requirement), the updated Standard Contractual Clauses of June 2021 (the mechanism that requires supplementary measures), and the EDPB recommendations on supplementary measures (the guidance on what those measures look like in practice). The architectural conversation that this article concerns is what follows from the legal framework: where the data must actually go, and what the manufacturer's architecture must look like to support it.

What data is actually being transferred

The first move in any cross-border data conversation is to identify what the data actually contains. This is harder than it sounds, because manufacturers tend to bundle data types together under the label "telemetry" and the GDPR analysis requires the bundle to be separated.

A plant generates several distinct categories of data. Operational telemetry from controllers — turbine speed, blade pitch, generator output, inverter status, battery state of charge — is, on its own, not personal data. It describes the equipment, not the people operating it. Condition monitoring data — vibration spectra, oil quality measurements, thermal profiles — falls into the same category. Aggregated performance data, plant-wide energy output, fault rates, mean time between failures — also not personal data.

Other categories are clearly personal data. CCTV footage from cameras in the control room, around the substation, in maintenance areas. Badge swipes and door access logs. Identifiable maintenance records — "Engineer X performed task Y on date Z" is personal data even if the task and the date are uncontroversial. The audit logs from the identity and access management platform are personal data by definition.

Some categories are mixed. A diagnostic data export prepared during a remote support session may include both the operational data being analysed and the screen recording of the engineer performing the analysis. A maintenance work order may include both the technical fault description and the credentials of the engineer who closed it. A condition monitoring report that names the inspector is personal data; the same report anonymised is not.

The exercise the lender's data protection function expects is data classification. Every data flow leaving the plant, or potentially leaving the plant, is identified, categorised, and assessed for personal data content. The result is a data inventory — a list of every data type, its source, its destination, its retention period, and its classification under the GDPR.

For most non-EU manufacturers, this exercise produces surprises. Diagnostic data flowing to the support cloud turns out to include identifiable engineer actions. Condition monitoring data turns out to in-

clude some plant operations metadata identifiable to the operations team. Telemetry channels that the engineering team assumed were purely operational turn out to carry contextual metadata that under GDPR analysis constitutes personal data.

The transfer impact assessment in practice

For each personal data flow that crosses the EEA boundary, the controller — the EU-headquartered project sponsor, in the archetype of this series — must conduct a transfer impact assessment under the framework established by Schrems II.

The assessment has a specific shape. Identify the destination country and the recipient legal entity. Assess that country's surveillance and intelligence laws — what access the authorities have to data held by entities in that jurisdiction, under what process, with what oversight. Determine whether the specific recipient is subject to those laws (most are, by virtue of being established in the jurisdiction). Determine whether the supplementary measures available — encryption with keys held only in the EEA, pseudonymisation that prevents re-identification, split processing across multiple jurisdictions — can close the gap identified.

For some destination countries the assessment concludes easily. Adequacy decisions exist for the United Kingdom, Switzerland, Japan, the Republic of Korea, New Zealand, Canada (commercial organisations), Israel, Argentina, Uruguay, the Faroe Islands, Guernsey, the Isle of Man, Jersey, Andorra, and — under the EU–US Data Privacy Framework adopted 10 July 2023 — the United States for recipients on the DPF list with valid certification; transfers to these jurisdictions proceed under Article 45 without further mechanisms. For some other countries the assessment is workable with supplementary measures — adequacy does not exist, but the country's surveillance framework

is bounded enough that encryption with EEA-held keys, or pseudonymisation, closes the gap.

For several major industrial manufacturing jurisdictions, the assessment is not workable in any practical configuration. The countries' national security and intelligence laws compel entities established in their jurisdiction to provide access to data they hold, under conditions and with oversight that fall short of what the EU considers necessary and proportionate. No technical supplementary measure short of refusing the transfer entirely satisfies the Schrems II test in these cases, because the entity holding the data is legally compelled to provide access regardless of what the encryption arrangements look like.

The lender's compliance team is not making a political judgement about these countries when they reach this conclusion. They are applying a legal test that the European Court of Justice has set out and that EU data protection authorities have elaborated. The conclusion is a matter of law. The remedy is architectural.

What the asset owner and lender will accept

Three architectural patterns survive the analysis.

EU-only data residency. All telemetry, all analytics, all support tooling, all backup and disaster recovery — terminates in data centres physically located in the European Economic Area, operated by legal entities subject to EU law, with encryption keys held by EEA-resident entities. Data does not leave the EEA at any point in the lifecycle. This pattern is the safest from the lender's perspective and is the one most major non-EU manufacturers have moved to for their European customer base. It typically requires the manufacturer to establish or contract with an EU-based subsidiary, EU-based cloud capacity, and EU-based engineering teams for any function that touches personal data.

In-host-country data residency. Data stays in the country where the plant is physically located — Egypt, in the archetype — under that country's data protection law, with appropriate safeguards. This pattern works when the host country has its own credible data protection framework and the EU lender is comfortable with the local regime as a sufficient floor. It is increasingly common for projects in North Africa, the Gulf states, and Latin America where local data residency requirements often align with the lender's preferences.

Hybrid residency with strict classification. Operational data with no personal data content flows to the manufacturer's home jurisdiction for analytics and product improvement purposes; personal data and identifiable data stays in the EEA or the host country under the patterns above. The hybrid model works only with rigorous data classification at the source, technical controls that prevent personal data from leaking into the operational channel, and continuous monitoring to confirm the separation is maintained. It is more complex than the other two patterns and is typically chosen by manufacturers who have substantial analytics investment in their home jurisdiction that they are not yet ready to relocate.

Patterns that do not survive include any flow where personal data passes through a non-adequacy jurisdiction in transit, any architecture where the cloud provider's home government has compelled-access powers over the encryption keys, and any arrangement where the manufacturer asserts compliance without producing the data flow diagram and the transfer impact assessment that demonstrate it.

The data flow diagram

A specific deliverable that the lender's data protection function will request is a data flow diagram showing every data type, every source, every transit point, every destination, every legal entity involved. The

diagram is the artefact that exposes architectural assumptions, and it typically reveals the situations the prose discussion misses.

A typical first-cut diagram, drawn honestly, shows flows the manufacturer's team did not consciously design. Backup flows that mirror primary data to additional jurisdictions for resilience. Analytics flows where data is replicated into a second processing environment and recomputed against a different model. Support workflows where ticket data, screen recordings and engineer logs cross borders as a normal part of the service organisation operating across multiple time zones. Disaster recovery arrangements that activate in a different region than the primary. CDN paths for software downloads that transit through edge locations the architecture team had not formally considered.

Producing this diagram honestly takes work. Many manufacturers find that the diagram they think they have differs significantly from the architecture as built — flows discovered through diagram production that nobody had documented, jurisdictions touched in transit that nobody had counted, legal entities involved that the procurement team had not separately assessed.

The diagram has a second use beyond the lender's review. It is the artefact the manufacturer's own data protection function uses to remediate gaps. Once the flows are visible, the architectural changes required — repointing a backup, relocating an analytics workload, restricting a support tool's data access — become specific and tractable, rather than the abstract task of "ensure GDPR compliance" that is impossible to engineer against.

At proposal stage

A manufacturer's bid that addresses cross-border data flow explicitly — with a data flow diagram, a classification of which data is personal, a residency commitment for personal data, a description of the trans-

fer mechanism for operational data, a documented transfer impact assessment for any non-adequacy destination, and the legal entity structure that delivers the architecture — is a bid that has anticipated the conversation. The lender's data protection function reviews the materials, identifies any remaining concerns, and the conversation proceeds.

A bid that does not address cross-border data flow, or that gestures vaguely at "secure cloud infrastructure" and "compliance with applicable data protection regulations", signals a gap that the lender's compliance team will surface during due diligence. The gap may be closeable through architectural change — repointing telemetry endpoints, establishing EU residency for personal data, building the data classification at source. The gap may be more fundamental, requiring the manufacturer to stand up new infrastructure or new subsidiaries to support the data flows that EU regulation permits. The closer the gap is to the manufacturer's core analytics platform, the more substantial the work to close it.

The deeper observation is that data residency is the area where regulatory specifics most directly constrain architectural choices. The cryptographic baseline, the patch contract, the logging discipline, the identity model — all generalise across global markets, and a manufacturer adopting them sees benefits beyond EU compliance. Data residency is structured by EU regulatory framework in a way that does not generalise as cleanly. Manufacturers serving EU-financed projects typically end up with EU-hosted analytics infrastructure, separate from their home-country infrastructure, with clean architectural separation between the two. The investment is significant. For manufacturers who serve EU-financed projects in any volume, it is also unavoidable.

The next article picks up the topic that the data flow review surfaces alongside data protection: the sanctions and provenance disclosure

the lender's compliance team conducts against the bill of materials, the suppliers, and the beneficial ownership of the manufacturer's supply chain.

This article reflects the regulatory landscape at publication. The EU's framework for international data transfers continues to evolve, particularly around adequacy decisions, EU-US data transfer arrangements, and EDPB guidance on supplementary measures. Court of Justice case law continues to develop. Specific transactions should be reviewed by qualified legal counsel rather than against this article. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

The bill of materials the sanctions desk will read

15 May 2026 · 11 min read · #compliance #security #industrial #oem-eu-readiness

At preferred-bidder stage, the lender's compliance team sends the manufacturer a request. It is not adversarial in tone. It is, on the contrary, almost mundane — a list of documents the manufacturer is asked to provide before financial close. The list runs to about a page. It asks for the manufacturer's bill of materials with manufacturer attribution for each line item, the corporate structure to the level of beneficial ownership, the compliance certifications for export control, and the manufacturer's confirmation that none of the items in the bill of materials, and none of the sub-suppliers, manufacturing partners, software vendors or component suppliers behind them, appear on any of the sanctions lists the lender's home jurisdiction and the syndicate participants are bound by.

The manufacturer's commercial team forwards the request to procurement. Procurement forwards it to legal. Legal forwards it to engineering for the bill of materials. Engineering needs three weeks to compile the data — the bill of materials at the requested level of granularity exists internally but has not previously been shared at this depth with any single customer. The corporate structure document is held by the parent company in a different jurisdiction. The export control certifications cover the home market but have not been mapped against EU equivalents. The sanctions cross-check has never been performed at the line-item level.

This is the sanctions and provenance disclosure pack. It is the lender's compliance team's routine deliverable, and it is, for many non-EU manufacturers, the first encounter with the level of struc-

tured supply chain transparency that EU project finance now requires.

The principle behind the disclosure is straightforward. The lender, as an EU-regulated financial institution, cannot finance a project whose supply chain includes sanctioned entities or items subject to export control violations. The obligation is not delegable. The lender must satisfy themselves, with evidence, that the project as financed does not breach any of the sanctions regimes the lender or any syndicate participant is subject to. The evidence comes from the manufacturer, in structured form, before financial close. After financial close, the manufacturer's ongoing obligation is to maintain that disclosure as conditions change.

What sanctions regimes apply

Multiple regimes, often overlapping, frequently applied as a union rather than individually.

The European Union maintains its consolidated sanctions list under the Common Foreign and Security Policy, implemented through Council Regulations. The list is updated continuously, expanded in response to geopolitical events, and applies to all EU lenders, all EU-headquartered borrowers, and any entity operating within the EU's jurisdictional reach. The lender's compliance function refers to this list as a baseline; entities or items on the list cannot appear in the project supply chain without specific licensing exceptions that, for industrial projects, are rarely granted.

The United States maintains parallel regimes through the Office of Foreign Assets Control. The Specially Designated Nationals list, the sectoral sanctions identifications list, and the entity list maintained by the Department of Commerce all apply to US persons, dollar-denominated transactions, and through secondary sanctions to many transactions involving non-US parties. If any syndicate participant

has dollar exposure, the OFAC framework reaches the project even if no other US connection exists.

The United Kingdom, post-Brexit, maintains its own sanctions framework through the Office of Financial Sanctions Implementation, generally aligned with EU positions but separately administered and not always identical in scope.

Norway, Switzerland and other non-EU European jurisdictions maintain national sanctions frameworks that typically align with EU positions but require separate cross-reference. United Nations Security Council resolutions impose sanctions that EU member states and most jurisdictions implement domestically.

The lender's compliance team typically applies the union of all relevant regimes — the most restrictive standard wins. A component acceptable under EU sanctions but flagged under OFAC will fail the lender's check if any syndicate participant has dollar exposure. A supplier acceptable under one regime but flagged under another fails on the regime where the flag exists. The manufacturer's disclosure must satisfy all of them simultaneously.

Three categories of sanctions matter for industrial supply chains. Asset freezes against specific named persons and entities — the consolidated lists. Sectoral sanctions against entire industries in specific jurisdictions, such as restrictions on advanced semiconductor exports to certain countries. And dual-use export controls, which cover items that can be used for both civilian and military purposes and which have their own regulatory framework alongside the sanctions lists.

What dual-use export control adds

Regulation (EU) 2021/821 — the EU dual-use regulation — covers items listed in its Annex I that require licences for export from the European Union. The list is extensive and includes many categories relevant to industrial control systems: certain semiconductors above

defined performance thresholds, certain encryption technologies, certain communication equipment, certain test and measurement instruments, certain materials and metals.

Parallel frameworks exist outside the EU under different names. The United States operates the Export Administration Regulations with its Commerce Control List, administered by the Bureau of Industry and Security. The United Kingdom maintains the Strategic Export Control Lists. Most major manufacturing jurisdictions have their own equivalents, frequently aligned through the Wassenaar Arrangement on Export Controls for Conventional Arms and Dual-Use Goods and Technologies.

For industrial equipment, the typical dual-use concerns include advanced semiconductors used in controllers and signal processing, encryption capabilities above certain key length or algorithm strength thresholds, communications equipment that could be repurposed for surveillance or military use, and certain categories of test and measurement equipment. The classifications are technical and granular; a specific semiconductor at one performance level may be uncontrolled while the same family at a higher performance level requires a licence.

The disclosure expectation is that the manufacturer can attest, with documentation, that any dual-use items in their bill of materials have been properly licensed for export to the project country, that no ongoing licensing constraints affect the project, and that no items in the supply chain are subject to recently-imposed restrictions that would prevent ongoing supply. The last point matters because dual-use restrictions can be tightened mid-project, with retroactive consequences for projects relying on continuous component supply.

The complexity arises when a manufacturer's product incorporates components originating in a jurisdiction with restrictive export controls, destined for use in a jurisdiction subject to those controls. A

controller assembled in one country, using a semiconductor manufactured in a second, deployed in a third, sold by a manufacturer headquartered in a fourth, may need to satisfy export control requirements in all four jurisdictions simultaneously. Each jurisdiction has its own framework; the union is what the lender's compliance team applies.

What the disclosure pack contains

The minimum content at preferred-bidder stage is six structured artefacts.

A bill of materials with manufacturer attribution for every component, both hardware and software. The software bill of materials earlier in this series feeds directly into this pack; the hardware bill of materials follows the same discipline but at the component-supplier level rather than the library level. Every line item identifies the supplier; for each supplier, the legal entity name and jurisdiction of establishment are recorded.

A corporate structure document showing the manufacturer's parent company, sister companies, subsidiaries, joint ventures, and ultimate beneficial owners. Beneficial ownership is typically defined as any individual or entity owning more than 25 per cent directly or indirectly, though some jurisdictions apply lower thresholds. The EU's 2024 anti-money-laundering package — Regulation (EU) 2024/1624 and Directive (EU) 2024/1640, replacing the prior AMLD framework, with the Regulation applying in full from 10 July 2027 — and parallel frameworks across other jurisdictions require disclosure to the asset owner or lender on request, with the manufacturer's parent typically holding the most current version.

A sanctions cross-check confirmation. The manufacturer attests, with evidence of the screening process used, that none of the entities in the bill of materials or the corporate structure appears on the consoli-

dated sanctions lists the lender applies. Most manufacturers run this screening through commercial compliance tools — Refinitiv World-Check, Dow Jones Risk and Compliance, LexisNexis WorldCompliance, Moody's Bureau van Dijk Compliance Catalyst, and others — that maintain current sanctions list data and provide an audit trail of the screening process.

Dual-use export control certification. Documentation confirming that any dual-use items in the bill of materials have been properly licensed for the project, including the licence numbers, the issuing authorities, and the validity periods. For items not requiring licences, a documented basis for that classification — typically a screening conducted by the manufacturer's export control function against the relevant control list.

Component provenance for high-risk categories. Specific declarations for semiconductors above defined performance thresholds, for encryption technology in the cryptographic chain, for communications equipment, and for any items subject to recently-imposed restrictions in any jurisdiction the lender considers material. Provenance is traced as deeply as the manufacturer's own visibility allows; gaps are flagged honestly.

The manufacturer's own beneficial ownership disclosure. The ultimate beneficial owners of the manufacturer's corporate structure, traced through holding companies and sister entities to the natural persons or sovereign entities with ultimate control. For privately held manufacturers, this may be straightforward. For listed entities with widely-held shares, the disclosure typically covers any owner above the threshold and the regulated shareholder reporting under the manufacturer's stock exchange rules.

Manufacturers who have submitted to EU project finance before have this pack pre-prepared and update it for each new transaction. Manufacturers new to EU project finance produce it for the first time, often

discovering during preparation that some of the underlying data has not been compiled before and requires a discovery exercise of its own.

Maintaining the disclosure

Sanctions regimes change continuously. A supplier acceptable at financial close may be added to a sanctions list six months later. A jurisdiction's status may change as a result of geopolitical events. A new sectoral sanction may apply to an industry the manufacturer operates in. A previously-licensed dual-use item may have its licence revoked, suspended, or rescoped.

The manufacturer's ongoing obligation is to monitor sanctions list changes and notify the asset owner promptly when their supply chain is affected. This is typically a contractual obligation written into the supply agreement, with specific reporting timelines — frequently within 30 days of the change, sometimes faster if the change is material. The reporting obligation includes the change itself, the affected components or relationships, and the manufacturer's proposed remediation.

Remediation may be simple — substituting a component, switching to an alternative supplier, restructuring a sub-contract — or it may be substantial, requiring re-certification of the modified equipment under the manufacturer's quality system. The asset owner's contracts typically include step-in rights, hold-back provisions, or material adverse change clauses that activate if the manufacturer fails to remediate within agreed timelines.

Sanctions screening as an ongoing activity is well-supported by commercial tools, but the operational discipline of monitoring, assessing impact, and reporting requires a compliance function the manufacturer maintains continuously. Manufacturers without an established compliance function build one specifically for EU-financed business;

this is one of the operational changes that EU project finance forces and that, once made, applies across the manufacturer's global business.

At proposal stage

A manufacturer's bid that includes — or proposes to provide on request, with a documented timeline — the full disclosure pack signals that the manufacturer has submitted to EU project finance before and has the supply chain transparency that the lender's compliance team requires. The conversation that follows is about specific items rather than about the existence of the pack: dual-use items requiring fresh licensing for the project country, borderline supplier relationships requiring clarification, beneficial ownership disclosures requiring drill-down, recent regulatory changes affecting the supply chain.

A manufacturer's bid that does not address sanctions and provenance, or that asserts compliance without producing the documentation, signals a gap that the lender's compliance team will surface immediately. The gap may be closeable through standard disclosure work — three to six weeks of compilation and screening for a manufacturer with reasonable internal records. It may surface specific items that require remediation, in which case the work expands. It may, in some cases, reveal supply chain relationships that prevent the manufacturer from being financed in the EU-financed project at all, in which case the bid is withdrawn or restructured.

The deeper observation is that sanctions and provenance disclosure is the topic where the lender's compliance team has the most direct authority. The cybersecurity conversations are mediated through technical advisers and the operational team; the data residency conversations are mediated through the data protection function. The sanctions check is conducted by the compliance team itself, against documented lists, with limited room for interpretation. An entity is on

the list or it is not. A failed sanctions check is a hard fail; there is no negotiating around it, and no commercial discount that compensates for it.

The next article picks up the lifecycle question that sanctions and provenance partly address but do not fully resolve: the support period for the deployed equipment, which under the Cyber Resilience Act must be declared and matched to the asset's operational lifetime, and which under the long-term service agreement runs on a different cycle entirely.

This article reflects the sanctions and export control landscape at publication. The EU consolidated sanctions list, OFAC, UK OFSI and parallel jurisdictional frameworks are updated through their respective channels, often in response to geopolitical events; the dual-use control lists evolve alongside them. Specific transactions should be reviewed by qualified compliance counsel rather than against this article. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

The five-year service contract and the twenty-five-year support commitment

15 May 2026 · 12 min read · #compliance #security #industrial #oem-eu-readiness

The project's commercial team has finalised the long-term service agreement with the turbine manufacturer. Five years from commercial operation date, renewable thereafter, with availability guarantees, response times, parts inventory commitments, and an agreed schedule for routine maintenance. The contract runs to fifty pages. Both parties' legal teams have reviewed it. It is signed.

The asset owner's security architect is reading the contract for a different purpose. She is mapping it against the Cyber Resilience Act requirement that the manufacturer publicly declare the support period for the product — the window during which security patches will be provided free of charge, through the patch delivery contract earlier in this series. She finds the declared support period stated in the manufacturer's product documentation, separately from the service agreement. The product support period is fifteen years from initial release.

The wind farm is being commissioned in 2027. The turbines being installed today are based on a product line first released in 2022. The declared support period therefore ends in 2037 — ten years into a twenty-five-year asset life. The long-term service agreement covers years one to five of that life, renewable for additional terms. The lender's term sheet requires the asset to remain supportable for the full operational lifetime.

Three timelines, none aligned, none addressing the same question. The service agreement protects availability and ongoing maintenance. The product support period protects security patching. The asset's operational life is longer than both. The architecture for closing

the gaps is the subject of this article, because none of the timelines closes them alone.

The principle is the distinction between the manufacturer's two commitments. The service agreement is a commercial arrangement — what the manufacturer will do for the asset owner, on what terms, for what price, over what duration. The product support period is a regulatory obligation — what the manufacturer must do for the product, regardless of any specific commercial relationship, under the Cyber Resilience Act and the equivalent regimes emerging in other jurisdictions. The two operate on different timescales, are governed by different frameworks, and impose different obligations on the manufacturer. They are not substitutes, and treating either as if it covered the other's scope is a category error the asset owner cannot afford to make.

The two commitments and what each covers

The long-term service agreement is the well-established commercial instrument. Five years from commercial operation date is the typical initial term, extendable by mutual agreement on three- to five-year cycles thereafter. The agreement covers availability guarantees (typically expressed as annual capacity factor or equivalent), mean time to recover for various fault categories, response time commitments for technical engagement, parts inventory and replenishment obligations, scheduled and unscheduled maintenance, performance warranties, and the commercial terms — fixed fee, variable element, indexation, payment schedule — that match the work to revenue.

The agreement is bilateral. It can be renegotiated at renewal, transferred under defined conditions, terminated for cause, novated as part of corporate transactions on either side. The pricing reflects the manufacturer's expected cost of providing the service over the

agreed term, with margin. Both parties have negotiating positions; the result is a contract.

The product support period is something quite different. CRA Article 13, read alongside the implementing acts as they are issued through 2026 and 2027, requires the manufacturer to publicly declare the support period for the product — the window during which security updates will be provided. The declared period must reflect the expected lifetime and use of the product class. For industrial equipment with a designed operational lifetime measured in decades, the support period must reflect that lifetime; declaring a five-year support window for a turbine controller designed to operate for thirty years is not a defensible position under the regulation's intent, even if it is technically a declaration.

The product support period is unilateral. The manufacturer commits to the period publicly, at the point of placing the product on the EU market. The commitment runs to all customers of the product, not specifically to any one of them. It cannot be renegotiated downward at the asset owner's request; it does not require any particular commercial relationship to continue; and it imposes the obligation to provide free security updates regardless of whether the asset owner has an active service agreement with the manufacturer.

The two commitments are independent in design, related in practice, and frequently confused in early conversations. The CRA does not require the manufacturer to provide field service free of charge for thirty years — the service agreement is the right place to address field service. The CRA does require the manufacturer to provide security patches for the declared support period, regardless of whether a service agreement is in place. The two work together when both are operating; they diverge when one is not.

What the CRA's declared support period actually requires

The declared support period is the manufacturer's public commitment, stated at the time the product is placed on the EU market, regarding the duration of free security update provision. The Cyber Resilience Act does not specify a minimum period in absolute terms; it specifies that the period must be appropriate to the product's expected use and lifetime.

For consumer products with short use cycles, support periods of a few years may be adequate. For industrial equipment, the Commission's evident expectation — reflected in the implementing acts and in guidance from ENISA — is that the support period align with the operational lifetime of the equipment in its intended deployment context. An industrial controller designed for twenty-five-year deployment will be expected to declare a support period commensurate with that lifetime.

The obligations within the declared support period are specific. Security updates must be provided free of charge. They must be made available without undue delay following the manufacturer's awareness of an actively exploited vulnerability — the 24/72/14-hour cadence from CRA Article 14 applies regardless of whether any specific customer has an active service agreement. The patches must be delivered through the structured contract described earlier in this series — signed artefacts, release documentation, SBOM diff, rollback procedure — not as silent updates or as customer-specific deliveries.

What the support period does not include is feature updates, performance improvements, or non-security maintenance. Those remain commercial offerings, typically delivered through the service agreement. The line between security and non-security work is sometimes contested in practice — a firmware update may include both — but the principle is that the security component is free within the support

period and the feature component is whatever the commercial arrangement specifies.

A particular operational consequence worth noting. The CRA's support obligation applies to the deployed product version, not only to the current product line. A turbine controller released in 2022 and deployed in 2027 is subject to a fifteen-year support obligation calculated from its original release date, regardless of whether the manufacturer's current product offering has moved on to a newer model. This matters for industrial equipment because deployment cycles are long: a controller installed today may be based on a product line several years old by the time it enters service.

What happens when the service agreement is not renewed

The scenarios in which the two timelines diverge are operationally important.

The service agreement reaches the end of its initial term and is not renewed. The asset owner may have negotiated for renewal at unacceptable commercial terms; may have decided to engage a different service provider for cost or quality reasons; may have moved to in-house operations and maintenance for the equipment. The product support period continues regardless. Patches still arrive through the patch delivery contract, free of charge, on the manufacturer's release cadence. What changes is the deployment capability — the asset owner's own team, the new service provider, or a third-party operations and maintenance contractor performs the field work.

The service agreement is renewed under a different service provider. Independent operations and maintenance companies — Vestas Multibrand Services, GE Vernova Services, Siemens Gamesa Services, but also independent providers like UpWind Solutions, Deutsche Windtechnik for cross-OEM work in wind, and several others across solar and battery — increasingly service multiple manufacturers'

equipment under common contracts. The original manufacturer's support obligation for security patches continues; the operational service relationship transfers. This pattern is becoming more common as renewable energy projects mature and asset owners optimise their operations and maintenance economics across mixed fleets.

The manufacturer is acquired by another company. The original entity's CRA support obligation transfers to the acquiring entity as a matter of regulation. In practice, the integration of product lines, engineering organisations and patch infrastructure under the new corporate structure may take months or years; the asset owner's contractual position should anticipate this, with commitments that survive corporate transactions and remedies that activate if the acquiring entity fails to maintain the support obligations.

The manufacturer exits the market or ceases trading. This is the scenario the safety net specifically addresses. The product support obligation continues in principle but is no longer practically deliverable by the original obligor. The asset owner's position depends on what was provisioned in the original procurement — the escrow arrangements, the transition support commitments, the source code access that the next section describes.

The safety net: escrow, transition support, decommissioning

Source code escrow is the most consequential element of the safety net for industrial equipment. For critical firmware components — programmable logic controller code, converter control software, safety-instrumented system code, the bootloader and secure boot chain, the protocol implementations — source code is deposited with a neutral third-party escrow agent under a deed of escrow. Common escrow agents for industrial software include NCC Group's Escode, Iron Mountain's escrow services, and Lloyd's Register for higher-assurance arrangements.

The deed specifies release triggers. The triggers commonly include manufacturer insolvency or cessation of trade, the manufacturer's failure to provide security patches within agreed timelines, the manufacturer's refusal or inability to remediate a critical vulnerability, and the expiry of the declared support period without a renewal commitment. On release, the source code becomes available to the asset owner or a designated successor service provider, with documented build instructions, dependency lists, and the technical artefacts required to make the code maintainable.

Escrow is more than the source code itself. It typically includes the build environment specifications, the test suite, the documentation of the development practices that 62443-4-1 conformance requires, the cryptographic signing keys (or their reset procedure), and the supplier contacts for any third-party components incorporated in the firmware. Without these, source code alone is not enough to continue maintenance; with them, a competent engineering organisation can take over the support function.

Transition support obligations are the second element. The manufacturer commits to a defined period of transition assistance if the service agreement is not renewed or if the manufacturer is exiting the market — typically twelve to twenty-four months, with documented deliverables: training for the successor service organisation, documentation packages, spare parts inventory transfer, and access to engineering support for specific technical questions during the transition window. The transition support obligation is contractual rather than regulatory; the asset owner negotiates it at procurement and the manufacturer commits to it as a condition of the project.

Secure decommissioning is the third element, addressing the end of the asset's operational life. When the equipment is retired, the cryptographic material it holds — private keys for PKI integration, stored credentials, certificates — must be securely erased before disposal.

The asset owner's PKI revocation infrastructure must record the revocation of the device's certificates. The configuration data must be sanitised. Physical disposal must follow the relevant waste electrical and electronic equipment regulations in the host country and any applicable extraterritorial obligations.

Each of these — escrow, transition support, decommissioning — is a deliverable the asset owner provisions at procurement and exercises at specific lifecycle points. None of them comes for free; each is a cost the manufacturer's bid must include. The cost is modest in absolute terms but disproportionately important if the unlikely scenarios materialise.

At proposal stage

A manufacturer's bid that addresses the lifecycle question explicitly — that states the declared support period and shows it aligned with the asset's operational lifetime, that includes source code escrow arrangements with documented release triggers, that commits to transition support with specific deliverables and timelines, and that addresses secure decommissioning — is a bid that has anticipated the conversation. The lender's compliance team reviews the lifecycle pack alongside the other submissions; the conversation that follows is about specific terms rather than about whether the manufacturer has thought through the long horizon.

A bid that conflates the service agreement term with the regulatory support obligation, or that proposes only the LTSA without addressing what happens at year six and beyond, signals a gap. The gap is not unbridgeable. A declared support period can be stated separately. An escrow deed can be negotiated with one of the established agents in three to four weeks. Transition support commitments are standard contractual language. The work is procurable from established sup-

pliers; the manufacturer's task is to recognise that the work is required.

The deeper observation is that the manufacturer's product strategy and the asset owner's project strategy operate on different timescales. The manufacturer's product roadmap may include end-of-life for a product line within ten years; the asset's operational life is two to three times that. The CRA's declared support period formalises this tension and forces the manufacturer to think about lifecycle commitments differently than historical practice has required. Manufacturers who internalise this — declaring meaningful support periods, building escrow into their standard offer, treating transition support as a service rather than a one-off — are positioned for a market in which lifecycle commitments are a procurement criterion rather than a contractual footnote.

The next article picks up the people and procedures that span all of these lifecycle questions: the personnel certifications, background checks, training discipline, and insurance arrangements that the manufacturer's service organisation needs to maintain across the support period, regardless of which commercial vehicle is in place at any given time.

This article reflects the regulatory and commercial landscape at publication. The Cyber Resilience Act's implementing acts continue to be issued through 2026 and 2027 and may clarify or alter specific support period obligations. References to escrow agents and operations and maintenance providers are illustrative rather than endorsements; specific arrangements should be reviewed by qualified counsel. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

The people, the certifications, the insurance

15 May 2026 · 13 min read · #compliance #security #industrial #oem-eu-readiness

Pre-mobilisation review meeting, three weeks before the manufacturer's service team is scheduled to begin commissioning work. The asset owner's HSE and security teams are present. The manufacturer's project manager has prepared the standard mobilisation pack — medical certificates for the engineers travelling to site, copies of their passports and visas, evidence of safety training, occupational health clearance, the personnel deployment list.

The asset owner's security team has additional questions. Has each engineer on the list been subject to a background check under the asset owner's vetting standard? What evidence of cybersecurity competency does each hold — relevant certifications, training records, prior experience on similar equipment? What insurance does the manufacturer maintain that covers cybersecurity incidents specifically — not the general professional indemnity, but the cyber-specific coverage with stated limits and clearly defined named perils?

The manufacturer's project manager works through the list with the asset owner's team. About half of the requested information is available immediately. The medical and safety training is documented. The visas and passport records are in order. The background checks are flagged as not previously requested at this depth, but the manufacturer can run them through their compliance partner with two to three weeks of notice. The cybersecurity competency evidence is sparse — most of the engineers have manufacturer-internal training records but few hold the independent certifications the asset owner expects to see. The cyber insurance question requires the manufac-

turer to come back with the policy schedule rather than the certificate summary that was attached to the mobilisation pack.

The meeting ends with action items running on the manufacturer's side, not the asset owner's. The mobilisation date slips while the gaps are closed.

The principle behind this review is that the operational assurance for the project rests on three connected foundations: the people who do the work, the qualifications that verify they are competent to do it, and the insurance that responds when things go wrong despite the qualifications and competence. Each layer depends on the others. An audit trail linking the three is the evidence base that satisfies the lender's compliance function, the asset owner's risk management, and the regulator's expectations under NIS2.

This is the human and commercial assurance loop. The technical pieces in this series have addressed the cybersecurity architecture of the project — what the equipment must do, how it must be configured, how it integrates with the asset owner's infrastructure. This article addresses the layer that operates the architecture in practice: the engineers who configure, maintain and update the equipment; their competency to do so; and the financial backstop that responds when their work, despite their competence, results in an incident.

The people: vetting, named disclosure, subcontracting

Background checks. The asset owner's vetting standard for engineers with access to OT systems typically includes identity verification, right to work in the project country, criminal record check covering the previous five to ten years depending on jurisdiction, reference checks, and, for engineers with privileged access, sometimes a financial probity check. For projects classified as critical infrastructure under host-country law, additional vetting may apply — formal security

clearance in some jurisdictions, equivalent processes administered by national authorities in others.

The vetting standard varies by jurisdiction and by the asset owner's group policy. NIS2 essential entities are expected to apply appropriate due diligence to personnel with access to network and information systems; Article 21 paragraph 2(i) lists human resources security among the risk management measures essential entities must implement. Some EU member states have specific frameworks for vetting personnel in critical infrastructure roles — the BSI (Bundesamt für Sicherheit in der Informationstechnik) in Germany for cybersecurity-specific vetting, with constitutional-level clearances administered through the relevant federal authorities where required, similar functions in France through SGDSN, the National Protective Security Authority in the United Kingdom. The asset owner's vetting practice typically aligns with the host country framework and the group policy, with the higher standard prevailing where they differ.

Named personnel disclosure. The asset owner requires advance disclosure of which specific engineers will be performing work on site or remotely. "Manufacturer's service team" as a generic descriptor is not acceptable; named individuals are identified, with their roles, certifications and vetting status documented. The named engineers are the people who receive identities in the asset owner's identity and access management infrastructure described earlier in this series; the personnel disclosure feeds directly into the identity provisioning workflow.

Subcontracting. The manufacturer's organisation often uses subcontractors — specialised technical services, local technicians in the host country, agency engineers for surge capacity, third-party operations and maintenance providers for specific functions. Subcontracting requires written consent from the asset owner, with the subcontractor's personnel subject to the same vetting and certification standards as

the manufacturer's direct employees. Undisclosed subcontracting is a contractual breach that the asset owner's audit function will surface; the appropriate practice is for the manufacturer to maintain a disclosed roster of approved subcontractors and to request specific consent for any work performed outside that roster.

The personnel pool that emerges from these requirements is typically smaller and more stable than the manufacturer's general service organisation. A handful of engineers, or a few dozen for a larger project, vetted, certified, named, and maintained as the project's approved pool. Manufacturers who serve EU-financed projects in volume maintain such pools as a standing capability rather than building them for each project.

The certifications: competency evidence

What the asset owner expects to see, beyond the manufacturer's own internal training records, is independent evidence of cybersecurity competency held by the engineers performing security-critical work. The competency landscape is well-developed.

For OT-specific cybersecurity work, the most commonly recognised certifications are the IEC 62443 Cyber Security Expert (CSE) certificates issued by accredited bodies including TÜV SÜD, TÜV Rheinland, ISA, exida, and several others. The ISA/IEC 62443 certification path runs through fundamentals, specialist and expert levels, with the expert level being the typical expectation for engineers responsible for security-critical work on industrial control systems. GIAC, the certification body affiliated with the SANS Institute, offers two industrial control specific certifications — the Global Industrial Cyber Security Professional (GICSP) and the Response and Industrial Defense (GRID) certification — both well-recognised in the OT community.

For broader cybersecurity foundations, the Certified Information Systems Security Professional (CISSP) from ISC2 and the Certified Infor-

mation Security Manager (CISM) from ISACA are commonly held by security architects and managers. The general SANS/GIAC certifications — GIAC Security Essentials (GSEC), Certified Incident Handler (GCIH), Certified Forensic Analyst (GCFA) — appear in the engineering teams responsible for incident response and forensic analysis.

For specific role areas, additional certifications apply. Network engineering for OT environments has industry-specific certifications from Cisco, Juniper, and the major OT networking vendors. Industrial protocol expertise comes through vendor certifications and through specialist training programmes. Safety-instrumented system engineering has its own certification ladder under TÜV functional safety programmes, separate from but adjacent to the cybersecurity certifications.

Country-specific equivalents exist and are usually recognised alongside the international certifications. The Bundesamt für Sicherheit in der Informationstechnik in Germany maintains training schemes; the Agence nationale de la sécurité des systèmes d'information in France operates national certification programmes; the National Cyber Security Centre in the United Kingdom maintains a certified cyber professional scheme; equivalent national programmes exist across the EU and in major non-EU jurisdictions.

The asset owner does not expect every engineer in the pool to hold every certification. The expectation is that the engineering team as a whole has the competency mix appropriate to the work, with specific individuals identified as the technical authorities for specific areas — a lead architect with CISSP and 62443 CSE, a senior incident responder with GCIH and GRID, an industrial network specialist with GICSP and vendor-specific certifications. The certifications are evidence; the underlying competency is what they signal.

The insurance: cyber liability beyond the standard policy

Standard professional indemnity insurance covers errors and omissions in the manufacturer's professional services. Standard product liability insurance covers defects in the products themselves. Neither typically covers the specific risks of cybersecurity incidents arising from the manufacturer's products or services in the way the lender requires.

Cyber liability as a distinct coverage emerged in the early 2000s for IT companies and has evolved to address industrial contexts. The lender's specification for cyber coverage typically requires several named elements. Named cyber coverage, explicitly identified as a separate line item in the policy rather than bundled into general professional indemnity. Third-party damages, covering damages to the asset owner arising from cyber incidents traced to the manufacturer's equipment or services. Ransomware response, covering forensics costs, negotiation support and remediation expenses, with appropriate sub-limits and clearly defined trigger conditions. Regulatory fines coverage, where insurable under the applicable law of the jurisdictions in which the manufacturer operates (some jurisdictions do not permit insurance for regulatory fines, in which case the policy should explicitly note the exclusion). Business interruption coverage for the asset owner's losses during outages caused by cyber incidents. Network security and privacy liability covering breaches that affect personal data or the asset owner's network beyond the manufacturer's specific scope.

The "silent cyber" issue is worth specific mention. Lloyd's Market Bulletin Y5381, issued in August 2022, required Lloyd's syndicates to address cyber risk in all policies — affirming coverage explicitly where it was intended, excluding it explicitly where it was not. The Bulletin took effect for new policies from January 2023. Parallel guidance from other insurance regulators followed, with similar effect — no-

tably the Prudential Regulation Authority's consultations on cyber underwriting in the United Kingdom, and corresponding statements from EU national regulators — reinforcing that the silent-cyber tightening is now market-standard rather than Lloyd's-only. The result was a tightening of cyber coverage in non-cyber-specific policies (which previously sometimes carried "silent cyber" exposure that responded to incidents the underwriters had not specifically priced) and a sharpening of exclusions in cyber-specific policies, particularly around acts of war, state-sponsored activity, and infrastructure attacks attributed to nation-state actors.

For the manufacturer, the practical implication is that their existing professional indemnity policy may not adequately cover cyber-specific incidents, and that standalone cyber policies require careful review of exclusions. The war and state-sponsored exclusions in particular have become broad enough in some policies that incidents involving sophisticated threat actors may fall outside coverage entirely. The lender's adviser examines the policy schedule rather than the certificate summary to confirm what is actually covered, with particular attention to the exclusions section.

Stated limits. The lender specifies the minimum cyber liability limit based on the project's risk profile. For utility-scale renewable energy projects in EU-financed structures, the minimum is typically in the tens of millions of euros for third-party damages, with sub-limits for ransomware response, regulatory fines, business interruption, and the other named elements. The limits are negotiated between the manufacturer's broker and the underwriter; the asset owner's adviser confirms that the limits as bound meet the project specification.

The audit evidence that links them

The three layers — people, qualifications, insurance — work together when each has documented evidence and the evidence is linked through the project's record-keeping.

For each named engineer in the project pool, the documented record includes the vetting record (background check completion, dates of completion, scope of the check), the competency record (certifications held with issuing bodies, expiry dates, scope of each certification), the authorisation record (what the engineer is authorised to do on the project, against which equipment, under whose supervision), and the activity record (what the engineer actually does, generated through the audit trail from the identity infrastructure described earlier in this series and the SIEM discussed alongside). The audit evidence is the documentation that ties an engineer's vetting status to their certifications to their authorisations to their actual activity. Without that linkage, the asset owner cannot evidence — to a regulator, to a lender, to an auditor — that the work performed on the deployed equipment was performed by appropriately qualified, vetted, authorised personnel.

The manufacturer's contribution to the audit evidence runs throughout the contract life. Maintaining current records for each named engineer in the project pool. Notifying the asset owner promptly when records change — an engineer leaves the manufacturer, a certification expires, a vetting result needs to be updated, a new engineer joins the pool and needs to be onboarded. Cooperating with the asset owner's periodic audit of the personnel records. Providing the insurance certificate renewal each policy year, with the policy schedule attached when changes have been made to the coverage. Responding to specific audit requests from the asset owner's compliance function or the lender's adviser within the timelines agreed in the contract.

The audit evidence is rarely scrutinised in detail — most of the time, the records sit in the asset owner's compliance archives and are referenced only at periodic audit points. When an incident occurs, the audit evidence becomes central. The investigation traces what happened, when, by whom, under whose authority, with what supporting credentials and qualifications. An incident response where the audit trail is complete, the credentials are current, the certifications are documented and the insurance responds is an incident that resolves cleanly. An incident response where any of those layers is missing or out of date becomes an incident the lender's compliance function reports differently than the asset owner would prefer.

At proposal stage

A manufacturer's bid that addresses the human and commercial assurance layer — that proposes a named pool of vetted, certified engineers for the project, attaches their certification evidence, describes the manufacturer's cyber liability insurance with the policy schedule attached for the lender's review, and outlines the audit evidence the manufacturer will maintain throughout the contract — is a bid that has anticipated the conversation. The lender's compliance team and insurance adviser review the materials, identify specific items requiring clarification, and the conversation proceeds.

A bid that treats personnel as a deployment-stage detail to be resolved after contract signature, or that addresses insurance only through a certificate summary without policy schedules, signals a gap that the late-stage review will surface. The gap is closeable, but the work involves the manufacturer's HR function (for vetting and personnel records), the manufacturer's professional development organisation (for certification evidence), and the manufacturer's risk management and broker relationships (for insurance restructuring). Three separate workstreams, all reaching procurement at the same late stage, all needing to complete before mobilisation can proceed.

The deeper observation: the human and commercial assurance layer often arrives later in the procurement process than the technical pieces. By the time the lender's insurance adviser is reviewing the policy schedule, the technical conversations have largely concluded and the deal is close to signature. The personnel and insurance work, if not anticipated, becomes the late-stage bottleneck that can stall an otherwise-ready project. Manufacturers who address it during the bid, alongside the technical submissions, find that the assurance layer is operationally straightforward — well-developed certification programmes, established vetting providers, mature cyber insurance markets. The work is procurable from established suppliers; the manufacturer's task is to recognise that the work is required.

The next article is the closing synthesis of the series — a procurement matrix that maps every topic across the seventeen substantive pieces to the procurement gate at which the conversation should be raised, from the initial request for information through mid-life review. The matrix gives a procurement team the operational tool the prose articles have been building toward.

This article reflects the regulatory, certification and insurance landscape at publication. The Cyber Resilience Act's implementing acts continue to evolve through 2026 and 2027; certification programmes are revised periodically by their issuing bodies; cyber insurance market conditions shift as underwriters respond to claims experience. Named certification bodies, training programmes and insurance market references are illustrative rather than endorsements; specific arrangements should be reviewed by qualified counsel and brokers. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.

Where each of these conversations belongs in the procurement timeline

15 May 2026 · 10 min read · #compliance #security #industrial #oem-eu-readiness

The procurement professional reading this series will, by now, have a procurement criterion against every substantive piece. The technical articles describe what the manufacturer must deliver, the regulatory articles describe why, the operational articles describe how the deliverables become living parts of the project, and the commercial articles describe how the assurance layer holds everything together. What has not been provided in any single artefact, until this one, is when each conversation belongs in the procurement timeline.

This article provides that artefact. The matrix below maps the fifteen substantive topics from articles 2 through 16 against the eight procurement gates that a renewable energy project typically passes through — from the initial request for information through the mid-life review that bridges the asset's first half-life into its second. The cells indicate the action taken at each gate for each topic. The prose around the matrix explains the dependencies — why an item at factory acceptance presumes an item already committed at contract, why an item missed at request for information typically cannot be retrofitted without commercial pain.

This is the piece a procurement team will print and keep on the desk through bid evaluation. The rest of the series will be referred to once or twice during a specific procurement; the matrix will be open continuously.

How to read the matrix

Fifteen rows, one per substantive topic, in the order the articles appeared. Eight columns, mapping the procurement timeline.

Request for Information (RFI). The initial market sounding, before a formal specification is issued. The asset owner asks open questions about manufacturer capability; the manufacturer's response shapes the subsequent specification.

Request for Proposal (RFP). The formal technical specification issued to qualified bidders, with the requirements stated in sufficient detail for the manufacturer to respond.

Bid Evaluation (Eval). The technical and commercial review of the manufacturer's response, with gaps identified for negotiation or for clarification.

Contract Negotiation (Contract). The final commercial and technical terms are agreed, deliverables and timelines are documented, and the contract is signed.

Factory Acceptance Test (FAT). The equipment is tested at the manufacturer's facility against the agreed specification, with the asset owner's commissioning engineer present and signing the FAT certificate.

Site Acceptance Test (SAT). The equipment is tested as installed at the project site, with integration to the asset owner's infrastructure verified.

Commissioning and Operations (Ops). The asset enters commercial operation and the topic becomes an ongoing operational discipline.

Mid-life Review. The periodic deep-dive review that asset owners typically conduct at year five and year ten of the asset's operational life, where major topics are revisited and renegotiated where necessary.

Each cell carries one of seven labels.

Raise — the topic is introduced for the first time, with the manufacturer asked to confirm capability or describe their current position.

Specify — the asset owner's specification states the requirement formally, with detail sufficient for the manufacturer to respond.

Evaluate — the manufacturer's response is reviewed against the specification, with gaps identified for negotiation.

Commit — the contractual commitment is made, with deliverables and timelines documented.

Verify — the deliverable is tested or evidence is examined, typically as part of acceptance.

Maintain — the topic becomes an ongoing operational discipline, with regular reporting or evidence collection.

Review — the topic is the subject of a periodic deep review, with the option to renegotiate or restructure.

An em dash indicates that the topic is not actively engaged at that gate. The matrix shows the gates at which the topic is in motion; the absences are where it sits dormant.

The procurement gate matrix

Topic	RFI	RFP	Eval	Contract	FAT	SAT	Ops	Mid-life
Communication network ownership (§2)	Raise	Specify	Evaluate	Commit	—	Verify	Maintain	Review
Substation boundary (§3)	Raise	Specify	Evaluate	Commit	Verify	Verify	Maintain	Review
Remote access architecture (§4)	Raise	Specify	Evaluate	Commit	—	Verify	Maintain	Review
Out-of-band components (§5)	Raise	Specify	Evaluate	Commit	Verify	Verify	Maintain	—
62443 documentation (§6)	Raise	Specify	Evaluate	Commit	Verify	Verify	Maintain	Review
Vulnerability disclosure programme (§7)	Raise	Specify	Evaluate	Commit	—	—	Maintain	Review
Software bill of materials (§8)	Raise	Specify	Evaluate	Commit	Verify	Verify	Maintain	Review
Cryptographic baseline (§9)	Raise	Specify	Evaluate	Commit	Verify	Verify	Maintain	Review
Identity and access (§10)	Raise	Specify	Evaluate	Commit	—	Verify	Maintain	Review
Patch delivery contract (§11)	Raise	Specify	Evaluate	Commit	Verify	Verify	Maintain	Review
Logging and SOC integration (§12)	Raise	Specify	Evaluate	Commit	Verify	Verify	Maintain	Review
Cross-border data flow (§13)	Raise	Specify	Evaluate	Commit	—	Verify	Maintain	Review
Sanctions and provenance (§14)	—	Specify	Evaluate	Commit	—	—	Maintain	Review
Support period and lifecycle (§15)	Raise	Specify	Evaluate	Commit	—	—	Maintain	Review
Personnel and insurance (§16)	Raise	Specify	Evaluate	Commit	—	Verify	Maintain	Review

What the matrix surfaces

Three patterns emerge from reading the matrix as a whole.

The RFI-critical topics. Fourteen of the fifteen rows have substantive activity at the request for information stage. Five of those rows describe topics that genuinely cannot be retrofitted later without considerable com-

mercial pain, and these deserve specific attention from the procurement team. Out-of-band components must be specified out at product-variant level before factory acceptance, or the manufacturer is asked to retrofit at FAT or SAT, with the cost the article on cellular modems described in detail. The remote access architecture must be addressed at RFI because the manufacturer's commercial model presumes persistent connectivity; raising it later forces a renegotiation of the service economics that the manufacturer's commercial team may not be authorised to conclude. The declared support period must be addressed at RFI because the manufacturer's product roadmap is committed in multi-year cycles and cannot be adjusted in the procurement window. The 62443 capability question must be addressed at RFI because building the capability where it does not exist takes months and cannot be compressed into a procurement schedule. The vulnerability disclosure programme must be addressed at RFI because building a public PSIRT takes four to six weeks at minimum and the question of whether one exists is binary.

A manufacturer whose RFI response is satisfactory on these five items is a manufacturer the project can proceed with. A manufacturer whose RFI response reveals gaps on any of them is a manufacturer whose bid will require parallel work to remediate, often with schedule consequences that surface in the integrated project programme.

The only row without an entry at RFI is sanctions and provenance, which is conducted by the lender's compliance team after the bidder list is short-listed rather than during the early market sounding. The sanctions screening typically opens at the RFP stage when the disclosure pack is requested as part of the technical submission.

The contract-stage commitments. Every row has a Commit entry at the contract gate. This is the structural argument the series has been making throughout: cybersecurity items in EU-financed renewable energy

projects do not exist as informal expectations or best-effort intentions; they exist as contractual deliverables, with timelines, with acceptance criteria, with consequences for non-performance. The manufacturer's commercial team must approach the contract negotiation expecting to commit to specific cybersecurity terms, in the same way they approach the negotiation expecting to commit to specific availability terms or specific performance warranties.

The operational disciplines. Every row except one has a Maintain entry at the operations gate. This is the second structural argument: cybersecurity expectations are not procurement criteria that close at commercial operation. They are operational disciplines that the manufacturer's service organisation will live with for the duration of the long-term service agreement and, in many cases, beyond the service agreement into successor arrangements. The single exception in the matrix is the out-of-band components row at the mid-life column, where a cellular modem physically removed at FAT does not need to be removed again at year ten; the matrix records the discipline that holds, rather than the action that recurs.

The mid-life review column is the least populated in the matrix. Most topics, by the asset's mid-life, are operating disciplines rather than fresh procurement items. The reviews that do happen at mid-life touch the topics where time has moved on — the post-quantum cryptographic migration from the cryptographic baseline article's longer-term concern, the long-term service agreement renewal cycle from the lifecycle article 's recurring decision points, the cyber insurance market re-procurement from the personnel and insurance article's continuing maintenance, and the periodic 62443 re-baseline that mature asset owners conduct as part of their wider cybersecurity programme review. These three or four items account for most of the mid-life work that touches the topics in this series.

At proposal stage, one more time

The matrix has one structural implication that bears stating explicitly, and that the series has been building toward across all seventeen articles. The cheapest stage at which to address any cybersecurity item is the earliest one at which it has substantive activity. The most expensive stage is the latest.

This is the inverse of the typical procurement experience for non-cybersecurity topics, where commercial terms are often deferred to the latest possible stage to maximise negotiating leverage. For the cybersecurity items in the matrix, the discipline runs the other way. A manufacturer's bid that addresses items at request for information in the manner the matrix describes is a bid that costs the manufacturer relatively little — they have done the analysis, they have prepared the response, they have begun the internal work where gaps exist. A manufacturer who defers the same items to contract negotiation, or to factory acceptance, is a manufacturer who will incur substantially higher costs to remediate under schedule pressure, often with the commercial price already fixed.

The procurement team's role, with the matrix in hand, is to surface the right items at the right gate. The technical team's role is to provide the evaluation criteria for each item at each gate. The asset owner's role, across both functions, is to make the discipline visible — to ask the questions early, to make the manufacturer's response part of the evaluation, and to refuse to let the right conversation happen at the wrong stage.

Closing the series

This is the seventeenth and closing piece of the series, following the anchor article. The series began with a scene of a manufacturer bidding into a renewable energy project in North Africa, surprised to find European

cybersecurity expectations attached to a project that physically sits outside European territory. The scene was specific because the conversation was specific. The series has worked through what the conversation actually contains — the regulatory framework that creates the expectation, the architectural disciplines that operationalise it, the documentary artefacts that evidence it, the lifecycle commitments that maintain it across the decades the asset will operate.

The substantive intent of the series has been straightforward. EU cybersecurity expectations in renewable energy projects are not as forbidding as they sometimes appear in the first conversation between asset owner and prospective manufacturer. They are extensive, but they are well-defined. The deliverables are mostly procurable from established suppliers. The architectural disciplines are mostly mainstream rather than exotic. What separates manufacturers who deliver into these projects routinely from manufacturers who struggle is not technical capability — both groups typically have it — but the procurement discipline of recognising what is required and engaging with it at the right stage.

The series is offered as a translation. Not of the law itself, which has its own languages, but of what the law lands as in the lender's term sheet, in the asset owner's specification, in the consultant's review, in the manufacturer's bid response. Translated into the operational vocabulary that engineers and commercial teams work in, the requirements become manageable. Manageable, in the framing the series has held to throughout, is most of what is required.

The series closes here. The articles remain available for reference; the matrix above remains the working tool; the prose above the matrix remains the rationale for any specific case where the matrix's instruction is contested.

The conversation that started in a meeting room in 2026, where a manufacturer's lead engineer was first asked questions they did not expect, can now happen in a different room, with a different bid pack, with a different vendor better prepared.

This article closes a seventeen-piece series. The procurement gate matrix reflects current practice at publication; gate names, sequence and emphasis vary between asset owners and between project structures, and the matrix should be adapted to the specific procurement framework of each project. Specific arrangements should be reviewed by qualified counsel and advisers rather than against this article. If a citation has rotted or a clause has moved, LinkedIn is the way to flag it.