
ASSURANCE YOU CAN EVIDENCE

A working manual for verifiable cybersecurity trust,
drawn from the AI Security Forum 2026

TABLE OF CONTENTS

Executive Summary.....3

1. From Labels to Measurable Trust.....5

2. AL1 — Evidence and Operational Compliance.....7

3. AL2a — Sovereignty Through Auditable Controls.....9

4. AL2b — Trust That Travels.....11

5. AL3 — From Compliance to Capability.....12

6. AL4 — Total Transparency and Continuous Verification.....13

7. The Proposed Assurance Ladder.....15

8. Testing the Ladder Against CSA 2.0.....16

9. What to Do Next.....18

Conclusion.....20

Appendix A — Session Context & Attribution Methodology.....21

 A.1 Event Overview & Structure.....21

 A.2 De-identification & Attribution Governance.....21

Appendix B — Supporting Reference Tables.....22

 Table 1 — Standards Mapped to the Assurance Argument.....22

 Table 2 — The Role of ENISA Across the EU Cybersecurity Framework.....22

EXECUTIVE SUMMARY

When an organisation must decide whether to trust a cybersecurity product or supplier, it has two options. It can examine the evidence: what was tested, by whom, against which requirements, and what can be independently inspected afterwards. Alternatively, it can focus on where the supplier is based and treat some countries of origin as inherently untrustworthy.

Based on the evidence provided by contributors at the ASF 2026 forum, taken together with wider industry perspectives gathered by Global Digital Foundation, this manual supports the first approach; a risk-based, evidence centric approach to assurance. Conversely, excluding suppliers because of their national origin is neither a stronger nor simpler form of assurance; it is a separate policy instrument. It produces no test result, audit trail, capability assessment or inspection finding. When used in place of technical assurance rather than alongside it as a distinct tool, it does not make an ICT ecosystem safer. Crucially flawed, it presents the absence of evidence as security.

Five findings from the morning track support this argument:

1. **Most security incidents do not start with an external attacker.** Data presented by an audit and legal practice showed that technical weaknesses and governance gaps together account for about seventy per cent of incident drivers, with external threats accounting for the rest. Two of the industry's largest recent outages involved no attacker at all, so screening suppliers by country of origin would not have prevented either one.
2. **Sovereignty can be delivered through control rather than exclusion.** France's SecNumCloud scheme — the EU's most demanding national cloud-security certification — meets its sovereignty goals through nineteen thematic domains of testable technical requirements. Its test is clear: control matters more than country of origin.
3. **Only evidence can travel across borders.** When results are portable, verifiable, comparable across schemes and independent of the testing laboratory, a product can be tested once and trusted in many markets. Country-of-origin screening, by contrast, produces no claim that can be tested.
4. **The main constraint on assurance is organisational capability, not the existence of a rule.** A regulator overseeing nearly three thousand entities said the practical problem is not a lack of requirements, but systems and controls that exist on paper without being used. Maturity — whether a control is embedded in daily operations — is what turns regulation into real security.
5. **Deep engineering transparency is possible and has worked in production, but it is costly.** One presentation described a supplier-funded evaluation centre that was governed independently, with priorities set by government rather than commercial interests. It was not withdrawn because it failed technically; its challenge was economic.

Building on these findings, this manual proposes a four-level “assurance ladder” — a way of organising evidence about a supplier or product, so that procurement conversations move from “where is this from” to “what can you show me”:

- **AL1** — documented and operational compliance
- **AL2** — verified technical control
- **AL3** — maturity-assessed capability
- **AL4** — engineering-transparent assurance

Two clarifications are important at this point. Firstly, this assurance ladder was not proposed at the forum; it is a synthesis, offered as a considered contribution by Global Digital Foundation. Secondly, this manual does not seek to dismiss sovereignty concerns, which are treated throughout as legitimate. However, we argue that strong assurance is best delivered through evidence-based approaches focussed on technical risk, matching controls to risk levels and application. Similar evidence-based approaches are advanced in other schemes such as EU Cloud Services (EUCS), or the French SecNumCloud scheme. Decisions about a supplier’s origin and decisions about technical assurance answer different questions, and neither can substitute for the other.

1. FROM LABELS TO MEASURABLE TRUST

There are two ways to answer the question “*can this be trusted?*” One asks where the supplier comes from. The other asks what was tested, by whom, against what requirement, and what can be inspected afterwards.

The first approach is fast, easy to communicate, and geopolitically convenient. The second is slower, more technical, and much harder to fake. The morning track at ASF 2026 was built around the second approach — not as a rebuttal of the first, but because its speakers work in the business of producing and relying on technical evidence.

AI Security Forum 2026 Quote: “Assurance is defined by what is tested, not by who makes the product.”

Assurance should be defined by what has been tested, not by who made the product. This quote came from an accredited evaluation laboratory operating across three continents, not from a supplier — which is itself telling: the case for evidence-based trust is being made by the bodies that produce the evidence, not by the companies being assessed.

Another speaker described the alternative just as precisely: an “*Illusion of Security*” built on assumption, emotion, bias and trust without evidence — memorably characterised as a “*colour of the flag*” approach. By contrast, “*Real Security*” rests on facts, assessment, transparency, engagement and collaboration. The distinction is practical as well as rhetorical: trust without evidence is easy to assert and maintain, but its real cost appears only when something fails.

AI Security Forum 2026 Quote: “Trust and resilience demand the capability to inspect each layer. Absent it, origin becomes the proxy for assurance.”

This manual treats origin-based screening not primarily as an ideological choice, but as a fallback used when an organisation or regulator lacks the capability to inspect a product directly. If that is right, the answer is not more argument but more capability: laboratories, certification schemes, transparency mechanisms and the trained people needed to operate them.

This capability fits within a broader regulatory landscape. European cybersecurity regulation is not a single instrument, but a set of interlocking frameworks: the EU Common Criteria scheme (EUCC) for technology-neutral, evidence-based assurance¹; the Cyber Resilience Act (CRA)² for products with digital elements; the Cybersecurity Act (CSA)³ and its certification schemes; the NIS2 Directive⁴, which establishes baseline risk-management and governance duties; sector-specific regimes such as DORA⁵ for financial services; and international standards including ISO/IEC 27001⁶ and the NIST Cybersecurity Framework⁷. Broadly, NIS2 defines the legal obligations, while certification schemes and standards provide recognised ways to show that those obligations have been met. [Figure 1](#) shows how these instruments connect.

Sections 2 to 6 show how evidence-based assurance works in practice, moving through four progressively deeper levels.

¹ EUCC, European Cybersecurity Scheme for Common Criteria https://certification.enisa.europa.eu/browse-topic/eucc_en

² EU Cyber Resilience Act <https://digital-strategy.ec.europa.eu/en/policies/cyber-resilience-act>

³ EU Cybersecurity Act <https://digital-strategy.ec.europa.eu/en/policies/cybersecurity-act>

⁴ NIS2 Directive <https://digital-strategy.ec.europa.eu/en/policies/nis2-directive>

⁵ Digital Operational Resilience Act https://www.eiopa.europa.eu/digital-operational-resilience-act-dora_en

⁶ ISO/IEC 27001:2022 Information security management systems <https://www.iso.org/standard/27001>

⁷ NIST Cybersecurity Framework <https://www.nist.gov/cyberframework>

FIGURE 1

The core European cybersecurity framework

Duties, evidence and methods form one assurance system




Figure 1 — The core European cybersecurity regulation framework: how EUCC, the CRA, the CSA, NIS2, sector rules and international standards interlock.

FIGURE 2

From proxy to inspectable evidence

Two decision routes - only one produces a testable claim




Figure 2 — From origin proxy to inspectable evidence: only the evidence route produces a testable assurance claim.

2. AL1 — EVIDENCE AND OPERATIONAL COMPLIANCE

Where do incidents actually start? The forum evidence points mostly to internal causes, not to who supplied the technology.

An audit and legal practice shared a breakdown of incident causes: 35% technical weaknesses, 35% organisational governance gaps, and 30% external threats. Taken together, the two internal categories account for roughly seventy per cent of what actually goes wrong.

FIGURE 3

Where incidents begin

Internal weaknesses and governance gaps dominate the evidence

70%
STARTS INSIDE




Figure 3 — Where incidents begin: technical weaknesses and governance gaps account for 70%, compared with 30% external threats.

Two case studies illustrated the point starkly, and neither was a breach in the conventional sense. A 2025 outage at a major content-delivery provider began with a routine database permission change that cascaded into a wider infrastructure failure, taking millions of services offline — with no attacker involved. A 2024 outage caused by a flaw in endpoint-security software crashed roughly 8.5 million systems worldwide, again with no attacker involved.

Both organisations in these examples were themselves top-tier security providers. No amount of supplier-origin screening would have prevented either event — the first, and largest, hole in the case for treating origin as a proxy for security.

This is the problem NIS2 is meant to address. It is not a defence against a particular attacker, but a governance framework requiring systematic risk management, clear accountability and strict incident reporting across eighteen essential and important sectors: an initial alert within 24 hours, a fuller report within 72 hours and a final report within one month. Its main pressure point is speed. Detection, supplier coordination and automated hand-offs between teams must now happen faster than the reporting clock allows, while organisations must also notify affected clients before regulators. However, Real-time supply-chain coordination remains undefined in the current framework.

Asked for a single piece of advice for this manual, one contributor put it bluntly: “***do not confine compliance to paperwork; build the process for real.***”

AL1 starts with evidence that governance operates in practice every day. It is much more than a compliance certificate to be filed away. Most organisations fail at this first level, and no higher level of assurance can make up for that gap.

3. AL2A — SOVEREIGNTY THROUGH AUDITABLE CONTROLS

The hardest question raised in the morning session was this: what should happen when a legitimate national sovereignty requirement meets a technical assurance framework? Does one have to give way to the other?

The answer on offer at the forum was not a general principle but a specific, working scheme. France's SecNumCloud⁸ v3.2 certification comprises nineteen thematic domains, built on an ISO/IEC 27001 base. SecNumCloud is based around six pillars, each written as an auditable requirement, see [table 1 below](#).

France's CNIL and Cigref organisations warn that replacing risk-based assessments with high-risk supplier assessments risks weakening France's gold standard SecNumCloud scheme⁹. Although primarily a privacy regulator, CNIL¹⁰ has a significant cybersecurity interest because GDPR requires organisations to implement "appropriate technical and organisational measures". Cigref¹¹ represents the interests of major French organisations that use digital technologies, including banks, insurers, energy and transport companies and public-sector bodies.

Table 1 – Six Pillars in SecNumCloud

Pillar	Requirement
Data sovereignty	All data on EU soil; administration and supervision from within the EU (Req. 19.2.c)
Ownership control	Non-EU entities capped at 24% individually / 39% collectively, direct or indirect (Req. 19.6.b)
Operational autonomy	Service maintainable via own competences or ≥2 independent suppliers (Req. 19.6.d)
Access control	Non-EU remote support only under real-time supervision via secure gateway, EU-based agent (Req. 9.7.d)
Legal immunity	No exposure to extraterritorial seizure without prior EU legal process
Third-party scrutiny	Full auditor scrutiny of uncertified components

AI Security Forum 2026 Quote: "NOT: Sovereignty by exclusion — BUT RATHER: Sovereignty through control. The test is not nationality — it is control. A US-origin provider can qualify IF it operates through an EU-majority-owned entity, under EU law, with EU-based staff and audited technical independence."

⁸ ANSSI SecNumCloud <https://cyber.gouv.fr/enjeux-technologiques/cloud/>

⁹ Article 'EUCS, the European cloud certification that threatens to disarm SecNumCloud' <https://www.cio-online.com/actualites/lire-eucs-la-certification-cloud-europeenne-qui-menace-de-desarmer-secnumcloud-15856.html>

¹⁰ CNIL <https://www.cnil.fr/en/cnil/cnils-missions>

¹¹ Cigref <https://digital-skills-jobs.europa.eu/en/community/organisations/club-informatique-des-grandes-entreprises-francaises-cigref>

FIGURE 4

Sovereignty through control

Six auditable pillars replace nationality with demonstrable control



Figure 4 — Sovereignty through auditable control: the six SecNumCloud pillars represented in Table 1.

The underlying point is that SecNumCloud demonstrates compliance through an evidence-based approach. The most demanding national cloud-security certification in the EU meets every sovereignty objective it was designed for without a single country-of-origin exclusion. Every requirement in it can be audited, contested and evidenced — and that is precisely what separates a genuine technical control from a label.

A control-based approach to sovereignty is not a novel idea. Telecommunications operators have worked under strict national sovereignty rules for decades — covering spectrum licensing, lawful interception, data retention and mandatory local presence — without anyone describing that regime as anti-competitive. Control-based sovereignty is, in other words, the normal operating practice of an entire adjacent industry.

One further figure reinforces the point: public administration accounts for 38.2% of all EU cyber incidents, making it both the sector with the highest concentration of threats and the one where certification is already contractually required.

Taken together, AL2 means that requirements have been independently tested against the actual product or service in question — not merely declared.

4. AL2B — TRUST THAT TRAVELS

AL2a shows that a control can be audited. AL2 also has a separate dimension: whether that audited control is recognised beyond its original context. Real evidence differs from a label in four ways, all of which must be present together: it is portable, so it can cross borders without retesting; verifiable, so it can be tested and disproved; comparable, so it can be assessed like-for-like across schemes; and laboratory-independent, so it remains valid even if the testing laboratory changes.

FIGURE 5

Trust that travels

Recognition depends on four properties working together




Figure 5 — Trust that travels: portable, verifiable, comparable and laboratory-independent evidence.

Of the four, verifiability is the decisive factor. An assurance claim that cannot be falsified is not really an assurance claim at all — and origin, by its nature, cannot be tested and cannot be proven wrong.

This is not to dismiss sovereignty concerns. The same contributor that made this argument also acknowledged that origin-based measures respond to real concerns — jurisdiction, supply-chain control, and coercion — and that such measures appear across all major jurisdictions, not just one political bloc. The proposition this manual draws from that acknowledgement, and the one the whole report rests on, is this: sovereignty instruments are legitimate in their own right, but a sovereignty instrument cannot do a technical assurance instrument's job, because it produces no testable claim about the product itself. Treating the two as interchangeable weakens both.

That erosion is already measurable in practice. The CCRA (the Common Criteria Recognition Arrangement) has broad international reach but is capped at evaluation assurance level EAL2. The voluntary EUCC scheme reaches a higher assurance level (AVA_VAN.3–5) but has narrower international recognition. Both are built on the same ISO/IEC 15408 foundation — and the further the two schemes drift apart, the more the overall system of mutual trust erodes. The Cyber Resilience Act, meanwhile, is pulling evidence-based assurance further into mainstream EU product law.

Taken with AL2a, this is the bridge to the next level of the ladder: control establishes that a requirement has been tested; recognition determines how far that result can travel.

5. AL3 — FROM COMPLIANCE TO CAPABILITY

The forum's most practical contribution regarding assurance came from a national regulator describing implementation as it happens in practice, rather than as it is imagined in principle. The regulator in question supervises nearly 3,000 entities. National cybersecurity registration was required by 30 June 2024, risk-management measures had to be in place by 18 October 2024, and the first formal audit cycle had to take place by the end of June 2026, repeating every two years thereafter; at the date of the forum, more than 1,000 audits had already been completed. A parallel national IoT certification scheme runs alongside this, at Basic, Substantial or High assurance levels.

The same speaker was candid about an unresolved architectural problem: once the EU's Cyber Resilience Act becomes fully applicable, the role of this national scheme may become uncertain, and the overlap between the two could force revision, alignment or partial replacement of one by the other. A working national scheme and an incoming horizontal EU regime have not yet been reconciled with each other — a practical assurance problem in its own right, independent of any wider geopolitical debate about sovereignty.

The sharpest point came from a panel discussion: what matters most is not whether a control exists on paper, but whether it is actually in use. The panel's experience was that there are controls that are not being implemented. Leadership commitment, not paperwork, was identified as the real differentiator.

FIGURE 6

From compliance to capability

A control creates assurance only when it is alive in the organisation




Figure 6 — From compliance to capability: implementation moves from registration and controls to audit and recurring verification.

Two contributions from opposite ends of the profession. A professional-services audit practice and a national regulator arrived independently at the same conclusion. The binding constraint on assurance is not whether a requirement exists. It is whether that requirement is actually alive inside the organisation. A control that exists only on paper is, from an attacker's point of view, no control at all.

6. AL4 — TOTAL TRANSPARENCY AND CONTINUOUS VERIFICATION

The deepest, and most demanding, form of assurance discussed at the forum came with an unusually candid account of its real costs. The speaker opened by explicitly declining to oversell certification, describing it as “**an entry ticket, not real-world or total security**” — a static solution to a threat landscape that keeps moving. Standards, in this framing, set the bar; they do not run the day-to-day defence.

At scale, that distinction becomes a governance model, not just a slogan. The example was a supplier-funded engineering evaluation centre for a single national market: funded and resourced by the supplier, but operated independently, free from supplier interference and accountable to an oversight board representing government, the supplier and network operators. The centre reported to government, to the supplier’s headquarters — not its local representative office — and to the relevant network carrier. The oversight board verified its independence, while national government set its priorities. It could not engage in pre-sales activity, hold contracts or promise customers delivery dates, because it did not control its own schedule.

It was precisely these constraints — not the word “transparency” itself — that made the model credible. An evaluation centre that a supplier can sell, schedule, or use as a sales promise produces evidence that is worth nothing.

As presented at the forum, the model did not fail on its own terms: reduced sales volume in that market made it economically unsustainable, even though the government had acknowledged it as a gold standard and network operators expected to keep using the equipment concerned for close to another decade.

The term used for this approach is “Total Transparency”, and its scope is easy to misread. It does not mean transparency from the supplier alone. It means continuous, proactive engagement from the supplier, the government and network operators alike; long-term commitment from all three parties; and significant investment in security infrastructure to protect the intellectual property held at the evaluation centre. The same account also conceded that the model was operationally difficult to sustain, and was ultimately funded by sales volume in that one market.

FIGURE 7

Total Transparency

Independent evaluation depends on governance, participation and scale




Figure 7 — Total Transparency: the independently governed evaluation-centre model and its operating preconditions.

This manual retains that caveat deliberately. Total Transparency depends on real preconditions: a government prepared to invest in oversight, operators willing to participate, and a supplier with enough commercial volume to bear the cost. It is not a cost-free model that can be replicated anywhere; any framework that presents it that way misdescribes how the mechanism worked in practice.

One possible shift in the economics, raised at the forum as a genuine but unproven possibility, is that AI could help codify evaluation expertise, making a wider range of these skills affordable across more networks. This is presented here explicitly as potential, not as demonstrated capability.

7. THE PROPOSED ASSURANCE LADDER

Sections 2 to 6 above were not written to fit a pre-existing template. They follow the order in which four different kinds of evidence were presented over the course of one morning, by practitioners who had not coordinated with one another and who work in different parts of the wider assurance system.

Read together, however, they arrange themselves into a natural sequence. Each of the four answers a different question about the same underlying product or supplier:

1. Is the governance process actually running?
2. Has the specific requirement been independently tested?
3. Is the control alive inside the organisation, not just on paper?
4. Can an independent third party inspect the underlying engineering itself?

These are not interchangeable questions. Each one costs more to answer than the one before it, and each tells you something the previous one cannot.

Set out in that order, they form a ladder of four assurance levels, AL1 to AL4, see [Table 2 below](#).

Table 2 – The Assurance Level Ladder

Level	Working label	The question it answers	Evidence
AL1	Documented & Operational Compliance	Does governance actually run — and can you show it?	Incident causes predominantly internal; process, not paperwork
AL2	Verified Technical Control	Was the requirement independently tested against the artefact?	portable, verifiable, comparable, lab-independent evidence
AL3	Maturity-Assessed Capability	Is the control alive in the organisation?	systems not in use; maturity bridges regulation and reality
AL4	Engineering-Transparent Assurance	Can an independent party inspect the construction, continuously?	oversight-board governance; total transparency from all stakeholders

AL2 in particular carries a second, independent dimension that should not be collapsed into the first: the breadth of recognition. A control can be recognised nationally (as with a national cloud scheme or a national IoT scheme) or internationally (as with EUCC, the CCRA, or CRA-harmonised standards). Depth of evidence and breadth of recognition are two separate questions, and a scheme can score differently on each.

Four constraints apply to how this ladder should be used:

1. It is not a rating. It does not score organisations against one another.
2. A higher level does not replace a lower one. An AL4 engineering-evaluation centre bolted onto an estate that will fail on a basic misconfiguration provides no real assurance.
3. AL4 is not universally replicable. Section 6 sets out the preconditions that made it possible in the case described there.
4. Its purpose is to move the discussion from a product's origin to the depth of evidence available about it.

It is worth noting that the forum also included the presentation of a separate four-phase model of national governance maturity, moving from unstructured exposure through reactive response to embedded and finally integrated governance. This model measures something different: national governance maturity, rather than the depth of assurance evidence available for a specific system or supplier. The two models sit on different axes and are complementary.

8. TESTING THE LADDER AGAINST CSA 2.0

A framework is useful only if it can be tested against a live policy question; otherwise, it is merely decorative. Start with what existing standards require. A review of eight international standards — ISO 31000, ISO/IEC 27001, ISO/IEC 27036, ISO 28000, ISO/IEC 22301, ISO 9001, ISO/IEC 27035 and ISO/IEC 12792 — found that none prescribes vendor exclusion or country-of-origin restriction as a control. The full mapping appears in Appendix B, [Table 1](#).

If none of the established standards call for origin-based restriction, where do such restrictions actually come from? Two distinct routes coexist within current EU instruments. One is standards-led: requirements flow into European standards, which in turn produce conformity evidence. The other is a supplier-risk route: jurisdiction, ownership and control considerations lead to a formal designation, which in turn leads to a restriction. Figure 8 traces how this second, supplier-risk route has entered the regulatory timeline over time.

FIGURE 8

EU cybersecurity law and policy, 2016-2026

Technical controls remain continuous while supplier-risk measures emerge selectively

TECHNICAL / ORGANISATIONAL CONTROLS




Figure 8 — European cybersecurity law and policy, 2016–2026: continuous technical and organisational controls alongside selectively emerging supplier-risk and non-technical measures.

The clearest recent example is the proposed Cybersecurity Act 2.0 (CSA 2.0), put forward on 20 January 2026 as a recast of the existing Regulation (EU) 2019/881. It adds a new Title IV covering ICT supply-chain security. Under the draft, once the European Commission formally designates a country as posing cybersecurity concerns, suppliers established in, or controlled from, that country become high-risk by legal definition — with consequences including exclusion from EU certification, from public procurement of ICT for key assets, and from EU ICT funding, and a phase-out period of roughly 36 months for mobile infrastructure.

Read against the assurance ladder set out above, the problem with this approach is precise rather than merely political. An origin-based restriction operates as a designation, not as a test. It produces no AL1 evidence that governance is running, no AL2 test result against a specific requirement, no AL3 assessment of organisational capability, and no AL4 inspection of the underlying engineering. It is not a higher rung on the ladder — it is a different kind of instrument altogether, operating on a different axis. Worse, by diverting investment away from testing, inspection and response capability, it risks eroding the very capability whose absence made origin-based screening feel necessary as a proxy in the first place.

Two important boundaries apply to this argument. First, nothing here suggests that sovereignty or economic-security concerns are illegitimate — section 4 above addresses those concerns directly and treats them as genuine. Second, the two contributions drawn on in this section were independent analyses presented separately at the forum; the synthesis connecting them to the assurance ladder is this report's own, not theirs.

9. WHAT TO DO NEXT

The forum's closing session asked every one of its six panellists the same question: what single insight would you offer for this assurance manual? All six answered it, and that exchange is, in a real sense, the origin of this document. Their answers converged, from six different professional vantage points, on similar themes: process matters more than paperwork; dialogue matters more than working in silos; manufacturers need to actually know their own products; leadership commitment is what makes controls real; markets need pragmatic solutions; and capability needs to be built, not assumed. It is worth being precise about what this represents: convergence from independent perspectives, not a negotiated consensus.

The guidance below is proposed on the strength of that convergence, and of the evidence set out in sections 2 to 8 above.

FIGURE 9

From policy to ecosystem

A single directional path for action




Figure 9 — From policy to ecosystem: a directional path for turning the forum's recommendations into assurance capability.



Policymakers and regulators. Before adopting an origin-based instrument, be explicit about which testable security property it establishes that existing technical controls do not already cover. If the honest answer is that the necessary inspection capability does not yet exist, that is a capability gap, and should be funded as one rather than papered over with a designation. Where a national scheme and a horizontal EU regime overlap, that overlap should be resolved explicitly — practitioners are already having to navigate it without clear direction.



Buyers and procurement authorities. Ask suppliers for evidence at a stated assurance level, rather than simply asking for a certificate. Evidence should come from operational records, not from policy or procedural documentation. For example, incident or configuration management records, test results from third-party labs etc.



Vendors and suppliers. The clearest supplier-facing finding of the morning came from the laboratory perspective: manufacturers frequently cannot fully document their own products, and struggle with the discipline required to maintain a software bill of materials (SBOM). AL2 is not reachable without that discipline. Publish what can be independently verified, and be explicit about what cannot yet be verified.



Laboratories and assurance bodies. The binding constraint on the system as a whole is recognition, not testing capability. Where results are not comparable across different schemes, portability fails regardless of how deep the underlying testing was — and the growing divergence between the CCRA and EUCC shows that this erosion has already begun.



The ecosystem collectively. Whether AL4-level transparency becomes a general option or remains a special case depends on the affordability of deep evaluation. The possibility that AI could make this more affordable is a hypothesis worth testing, and it is presented here explicitly as a hypothesis, not as an established fact.

CONCLUSION

We have argued in this manual that strong assurance is best delivered through evidence-based approaches focussed on technical risk, matching controls to risk levels and application. This manual proposes a four-level “assurance ladder” as a way of organising evidence about a supplier or product, so that procurement conversations move on from “where is this from” to “what can you show me”. Real evidence is portable, so it can cross borders without retesting; verifiable, so it can be tested and disproved; comparable, so it can be assessed like-for-like across schemes; and laboratory-independent, so it remains valid even if the testing laboratory changes.

The forum participants, when discussing assurance, independently and from six quite different professional positions, supported an evidence-based approach and each rejected the shortcut of treating a supplier's country of origin as a substitute for actual evidence.

- The auditor explained that most incidents start on the inside, where a flag on a country of origin has no purchase at all.
- The certification assessor achieved their sovereignty goals by including a sovereignty objective into nineteen domains of auditable technical requirement.
- The laboratory observed that origin yields no claim that can actually be tested, and therefore none that can be trusted.
- The regulator pointed out that existing controls are frequently not actively implemented. The message here is that implementing controls and gathering evidence is more effective than labelling a supplier based on origin.
- The evaluation centre supported the evidence-based approach by publishing what an independent party could inspect, and by being transparent about who controlled its own schedule.
- The policy researcher pointed out that supplier restrictions actually originate in policy and law, not in the technical standards those policies invoke.

The common thread running through all six of these positions is this manual's central argument: where genuine technical evidence exists, a supplier's country of origin adds nothing that the evidence has not already settled. Where that evidence does not exist, origin is not a legitimate substitute for it — it is a symptom of the underlying gap in assurance capability.

The assurance ladder proposed in this manual is one way of organising that evidence. It is offered for challenge and refinement, not for immediate adoption. As one contributor put it in the forum's closing remarks, the harder discipline is being honest about what each individual instrument cannot do on its own.

APPENDIX A — SESSION CONTEXT & ATTRIBUTION METHODOLOGY

A.1 Event Overview & Structure

The insights in this manual are drawn directly from the morning track of the AI Security Forum (ASF) 2026, held at Sapanca Lake, Türkiye, in June 2026. The morning programme comprised six specialist keynotes and an executive panel covering the end-to-end security assurance ladder:

- Cybersecurity Incident Analysis: Empirical failure drivers and operational outages (audit and legal practice).
- National Certification in Practice: Sovereignty through auditable controls under SecNumCloud v3.2 (national certification assessor).
- International Standards & Certification: Core properties of portable, lab-independent evidence (international evaluation laboratory).
- National Security Frameworks: Regulatory supervision and maturity-bridged compliance (national supervisory authority).
- Transparency & Evaluation Centre Models: Deep engineering evaluation and Total Transparency (engineering evaluation centre).
- Cybersecurity Maturity & Supply Chain Policy: Governance models and proposed CSA 2.0 Title IV provisions (policy research institute).

A.2 De-identification & Attribution Governance

No Attribution by Default: to preserve candid technical exchange, no speaker or entity is named in the body text. All evidence is referenced by organizational role or accredited lab type.

Editorial Independence: the proposed Assurance Ladder (AL1–AL4) is a synthesis framework introduced by this report. It does not constitute an official endorsement, standards proposal, or binding consensus by any participant.

Evidence Grounding: all empirical data, statutory timelines, and architectural requirements strictly reflect speaker presentation decks and verified session transcripts from ASF 2026.

Reordering disclosure: the policy research institute's contribution was positioned sixth and last on-site, as the closing keynote. In this draft it has been moved to follow the Assurance Ladder so that the affirmative argument is established before the CSA 2.0 critique. The true on-site chronological order is preserved in the session record.

APPENDIX B — SUPPORTING REFERENCE TABLES

Table 1 — Standards Mapped to the Assurance Argument

Function	Standard	Description	Content
Steering	ISO 31000	Risk management	Evidence-based risk assessment & treatment across the full lifecycle
Authority	ISO/IEC 27001	Information security management	Auditable ISMS embeds security, privacy & continuous improvement
	ISO/IEC 27036	Supplier relationships	Security governance across multi-tier hardware / software / service supply chains
Continuity	ISO 28000	Security management system	Supply-chain risk identification & mitigation, end-to-end operational resilience
	ISO/IEC 22301	Business continuity & resilience	Prevent, withstand & recover from disruption, integrates cyber incidents
Learning	ISO 9001	Quality management	Process discipline & auditability, repeatable, verifiable practice
	ISO/IEC 27035	Incident management	Detect, respond, escalate & learn from cybersecurity incidents
Visibility	ISO/IEC 12792	AI transparency & explainability	Taxonomy for measurable explainability, auditability & comparative assurance

Table 2 — The Role of ENISA Across the EU Cybersecurity Framework

The European Union Agency for Cybersecurity (ENISA¹²) is the EU's specialist agency for cybersecurity. Established in 2004 and given a permanent mandate under the Cybersecurity Act (Regulation (EU) 2019/881), ENISA's role is to strengthen cybersecurity across the European Union by supporting Member States, EU institutions and industry.

Regulation	ENISA's role
Cybersecurity Act (CSA)	Central role. Hosts the European Cybersecurity Certification Framework and develops candidate certification schemes such as EUCC.
NIS2 Directive	Produces implementation guidance, supports the CSIRTs Network and EU-CyCLONe, develops incident reporting guidance and helps Member States achieve consistent implementation.
Cyber Resilience Act (CRA)	Provides technical expertise and vulnerability management guidance; expected to support implementation and maintain technical knowledge relevant to secure products with digital elements.
DORA	Supports financial-sector resilience through technical guidance, cyber exercises and cooperation with sector regulators where appropriate.
EUCC certification scheme	Technical development and maintenance of the scheme, working with Member States and certification bodies.

¹² ENISA <https://www.enisa.europa.eu/>