The Operational Technology Threat Landscape, 2026: From Access to Action

Bruno Salmazo
Liscere
May 2026 · Liscere Briefing · LB-2026-01 · v1

Executive Summary

Attacks on industrial systems are no longer rare or theoretical. Across 2025, survey and incident data describe a sector under sustained pressure, with operational disruption a common outcome. The pattern behind the totals matters more than the totals themselves. Adversaries increasingly reach industrial networks through legitimate access and then act like operators, issuing commands that are valid at the protocol level and consistent with normal use. Defences built to check whether traffic is well formed, whether a sender is authorised, or whether activity is statistically unusual do not answer the question that now decides the outcome: is this otherwise valid action appropriate for the state the plant is in? This briefing sets out the evidence for that shift and the gap it leaves.

1. Why this matters now

For most of their history, operational technology (OT) networks were isolated and built for availability rather than security. That isolation has eroded. As industrial and corporate networks converge, the assumption of a trusted perimeter no longer holds, and the consequences are measurable. In the 2025 SANS State of ICS/OT Security survey, more than one in five organisations reported a security incident in the preceding year, and around four in ten of those incidents disrupted operations [1]. Ransomware has scaled in parallel: Dragos tracked 119 ransomware groups affecting roughly 3,300 industrial organisations during 2025, up from 80 groups the year before, with manufacturing accounting for more than two thirds of victims [2]. At European level, the EU cybersecurity agency analysed thousands of incidents across its 2024 to 2025 reporting period and recorded escalating, coordinated targeting of critical infrastructure by both state-aligned groups and hacktivists [4]. The message across these independent sources is consistent: industrial cyber incidents are now frequent enough, and disruptive enough, to be treated as an operational reality rather than a tail risk.

2. From breaking in to acting like operators

The more consequential change is in how attackers behave once inside. Access itself is increasingly the way in rather than the obstacle: in the SANS data, around half of reported incidents stemmed from unauthorised external access, yet only a small minority of organisations had OT-aware access controls such as session recording or industrial authentication in place [1]. The Dragos year-in-review describes adversaries who do more than gain a foothold. They learn how a control system operates, where its commands originate, and how a sequence of valid operations can produce a physical effect [2]. Threat groups have been observed working from engineering workstations, the consoles operators use to change controller logic, and mapping the control loops that drive a process. Sophistication is not a prerequisite. When an activist group reached and manipulated programmable controllers at water utilities, the technical means were simple, the interactions were valid, and the disruption was real [3]. As an attacker's activity becomes indistinguishable from an operator's, the signal that something is wrong moves from the network to the process itself.

3. Why current defences leave a gap

Established OT defences each answer a necessary question, and none answers the one that now decides the outcome. Protocol checks confirm that a message is well formed and uses an expected industrial protocol; many field protocols, including Modbus, carry no authentication of their own, so a well formed command from a reachable station is simply executed [7, 5]. Access control confirms that a sender is permitted to act. Anomaly and intrusion detection flag traffic that is statistically unusual or matches a known pattern. A command can clear all three checks, being well formed, sent over an authorised channel, and unremarkable on the wire, and still be wrong for the moment: a configuration change that is appropriate during maintenance but not during production, or a setpoint that is valid in isolation but unsafe given the current process state. This is part of why so much detection still fires only after the process behaves abnormally [2], by which point the action has already taken effect. The missing layer does not ask whether an action is permitted. It asks whether it is appropriate.

4. The durable question

The structural trend behind the 2025 data is unlikely to reverse. Connectivity continues to grow, adversaries continue to learn the process, and the line between a malicious action and a legitimate one continues to blur. The defensive question that follows is durable. Given a protocol-valid, authorised action against a known industrial asset, in a known operational context, is that action appropriate, and what evidence supports the decision? This is the question the Liscere framework is built to evaluate. It treats an industrial action, rather than a packet or a user, as the unit of evaluation, and assesses it across the protocol operation, the affected automation artefact, the declared operational context, and the applicable policy constraint. The framework and an initial Modbus/TCP evaluation are documented separately [8, 9]; this briefing is concerned only with why that question has become difficult to avoid.

5. A European note

For operators in the European Union, the trend arrives alongside regulatory pressure. The EU cybersecurity agency reports sustained targeting of energy and water operators, including internet-accessible OT management interfaces, by hacktivist and state-aligned actors across the bloc [4]. At the same time, the NIS2 Directive and the IEC 62443 series raise expectations for monitoring, segmentation, and least privilege in industrial environments [6, 5]. These frameworks motivate context-aware control of individual actions, but they do not themselves implement it, leaving operators to close the distance between what is permitted and what is appropriate.

6. What to watch

Three developments are worth tracking. First, the continued movement of adversaries from reconnaissance toward deliberate manipulation of process behaviour, which raises the value of evaluating actions rather than only detecting intrusions. Second, the expansion of OT to distributed and remote assets, which widens the surface over which a valid-looking command can be issued. Third, regulatory consolidation, which is likely to make explainable, evidence-backed decisions about industrial actions a compliance expectation as well as a security one.

This briefing synthesises publicly reported data as of mid-2026 and reflects the view of the Liscere research initiative. The Liscere framework and its evaluation are reported in LTR-2026-02 and LTR-2026-01.

References

[1] SANS Institute. State of ICS/OT Security 2025. SANS Institute, 2025.

[2] Dragos. OT Cybersecurity Year in Review 2026 (reporting on 2025). Dragos, Inc., 2026.

[3] Dragos. OT Cybersecurity Year in Review 2025 (reporting on 2024). Dragos, Inc., 2025.

[4] European Union Agency for Cybersecurity (ENISA). ENISA Threat Landscape 2025. ENISA, 2025.

[5] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson. Guide to Operational Technology (OT) Security. NIST Special Publication 800-82, Revision 3, 2023.

[6] International Electrotechnical Commission. IEC 62443-3-3:2013, Industrial communication networks, Network and system security, Part 3-3: System security requirements and security levels. IEC, 2013.

[7] Modbus Organization. Modbus Application Protocol Specification V1.1b3, 2012.

[8] B. Salmazo. Operational Legitimacy of Industrial Control Actions: The Liscere Framework. Liscere Technical Report LTR-2026-02, 2026.

[9] B. Salmazo. Context-Aware Evaluation of Industrial Control Actions: A Modbus/TCP Baseline in the OT Lab. Liscere Technical Report LTR-2026-01, 2026.