Passive and Active in Operational Technology: A Field Guide to the Trade-offs

Bruno Salmazo

Liscere

July 2026 · Liscere Briefing · LB-2026-05 · v1

Executive Summary

Passive and active are often posed as rival methods, but in operational technology they are better understood as tools with different reaches, chosen against a calculation that differs fundamentally from the one made in IT. This briefing separates two questions that the debate usually conflates, discovery and inventory on one axis, continuous monitoring and detection on the other, and weighs the methods on each. The reason OT is not IT is concrete, not cultural: industrial devices are fragile, priorities do not tolerate unplanned outages, and equipment stays in service for decades. Against this backdrop the trade-offs are clear and symmetric. Active methods reach silent devices and recover metadata that traffic does not carry, at a disruption risk that modern OT-safe scanners reduce but do not eliminate. Passive methods introduce no traffic and observe continuously, at the cost of blindness to silent devices, dependence on a vantage point, and, increasingly, opacity under encryption. Federal and sector guidance reflects this by establishing passive monitoring as the baseline and placing the burden of proof on active methods. Encryption adds a further axis, on which hardware measurement shows that what an encrypted channel changes is less the method of observation than its location. The mature pattern is hybrid, and the question each defender is left with is whether the horizon of their chosen method reaches far enough for what their own environment asks of it.

1. Two axes, not one

Discussions of passive versus active in operational technology often collapse two distinct questions into one, and the resulting confusion obscures more than it clarifies. The two belong to different layers of a security programme and are answered by different tools.

The first axis is discovery and inventory: establishing what exists on the network, which devices are present, how they are configured, and what firmware they run. This is a periodic question, answered well by a snapshot taken at a point in time.

The second axis is monitoring and detection: observing the network continuously to understand how devices behave, what they communicate, and when something departs from the norm. This is a continuous question, answered by an ongoing vantage on traffic rather than a snapshot.

Both axes admit passive and active methods, but the trade-offs differ between them. An active discovery scan and an active monitoring probe carry different risks and yield different information; passive discovery and passive monitoring have different blind spots. Treating passive versus active as a single choice forces a decision that should be made twice, once per axis, against different criteria. The remainder of this briefing keeps the two separate, and gives more weight to the monitoring axis, where the continuous vantage on traffic is the defining constraint.

2. Why OT is not IT

The passive-versus-active calculation is routine in IT and fraught in OT, and the reason is not conservatism. It is that the devices, the priorities, and the timescales are genuinely different.

The devices are fragile. Many industrial controllers have constrained processing power and network stacks that were never hardened against unexpected input. Sending them traffic they did not solicit is not a neutral act. The effect is documented rather than hypothetical: in one systematic study of ICS asset-discovery tools, an ordinary Nmap scan against a programmable logic controller drove the device into a failure state, with status lights flashing and a full power cycle required to recover, a result the authors noted was reproducible by anyone using the tool in a standard way with no specialist skill [1]. Field reports describe specific models that reliably fault when interrogated by common IT scanning tools [2]. The vendor response in the first case, that the device was end-of-life, is itself instructive: the model remained widely deployed and available for sale. In IT, a scan that crashes a host is an inconvenience; in OT, the host may be controlling a physical process, and the failure propagates from the network into the world.

The priorities are inverted, but the shorthand for this is often overextended, and it is worth stating carefully. IT security conventionally orders its goals as confidentiality, integrity, and availability, and OT is frequently described as reversing this to availability first. The reversal is real enough that IEC 62443 treats the difference explicitly [3]. But the "availability first" ordering began as a corrective for IT practitioners entering OT, not as a security architecture principle, and the major OT attacks, Stuxnet, Triton, and Industroyer, targeted integrity, not merely availability [4]. Neither NIST SP 800-82 nor IEC 62443 endorses a fixed priority ordering of the three properties; both derive their requirements from threat modelling and zone classification rather than from a ranking [4]. The accurate statement is narrower and sufficient here: an OT environment does not tolerate an unplanned outage the way an office network does, and a method that risks causing one is judged against a far less forgiving standard.

The timescales are long. IT hardware turns over every three to five years; OT systems are designed to run for two decades or more [3]. A method must therefore work not against the current generation of devices but against generations that predate modern network hardening by many years, and that will remain in service long after they stop receiving updates. The fragility described above is not a transitional problem that replacement will solve; it is a standing property of a field whose installed base is measured in decades.

3. What active methods do well, and what they cost in OT

Active methods send traffic to devices and reason from the responses. Their strength is reach: they can obtain what passive observation cannot, because they do not wait for a device to reveal itself.

Three capabilities stand out. Active methods find silent devices, those that are present but generate little or no traffic of their own, and which a passive observer therefore never sees [5][6]. They recover device metadata, firmware version, model, and configuration detail, that is often absent from ordinary network traffic and can only be obtained by asking [6][7]. And they do not depend on a vantage point: where no tap or mirror port exists, an active method can still reach a device across the network, whereas a passive one is simply blind there [8].

The cost in OT is the risk described in the previous section, and it is not to be minimised. But a fair account must also record that active methods have evolved. The aggressive, IT-native scanning that gave the approach its reputation, malformed packets, security probes, and high packet rates against fragile hardware, is not the only form active discovery takes today. Purpose-built OT scanners constrain themselves: strictly standard-conforming traffic, no security probes, per-host rate limiting on the order of tens of packets per second, and communication in each device's native industrial protocol rather than generic port interrogation [9][6]. Some vendors extend this to controlled, read-only, protocol-aware queries scheduled within planned maintenance windows [5]. This is a genuine change, not a rebranding: the modern OT-safe active query is a different instrument from the Nmap sweep that crashes a controller.

Two qualifications keep this in proportion. First, "OT-safe" is a claim about typical devices and tested conditions, not a guarantee for every controller in every state; the fragility of the oldest hardware means some devices remain unsafe to query at all, and that residual set does not shrink to zero. Second, active

methods still miss what is offline or transient at the moment of the scan; a snapshot sees only what is present when it is taken [10]. Reach is real, but it is not completeness.

4. What passive methods do well, and where they are blind

Passive methods observe traffic without participating in it, typically from a span port or network tap. Their defining strength is the mirror image of the active method's defining cost: they introduce no traffic, and so introduce no risk of disrupting the devices or the process. On a network where an unsolicited packet can fault a controller, a method that sends none is safe by construction [5].

That safety comes with a second advantage suited to the monitoring axis: passive observation is continuous by nature. It does not sample the network at intervals; it watches the traffic as it flows, which is what detection of behaviour over time requires. And it needs no credentials and no access to the devices themselves, only a vantage on the traffic they already exchange.

The blind spots are equally definite, and pretending otherwise would weaken the account. A passive method sees only what talks: a device that is silent, dormant, or transmitting infrequently may never appear at all, and the risk it carries goes unobserved [6][7]. Metadata that the devices do not put on the wire, firmware and configuration detail, is not recovered by listening [7]. The method depends on a vantage point, and taps and span ports are not present on every segment, so coverage has gaps wherever the mirror does not reach [8]. And it needs time: a passive baseline accumulates only as traffic is observed, so the picture is partial until enough has flowed [8]. To these must be added the constraint examined in the companion briefing on S7 encryption: where traffic is encrypted, passive observation of the payload recovers only transport metadata, and the substance of the communication is out of reach [14].

The pattern across this section and the last is not that one method is sound and the other flawed. It is that each method's strength is the other's blind spot, and that in OT the axis of difference that matters most, whether a method can disturb the process, falls decisively on one side.

5. What the standards say

The preference that OT practice shows for passive methods is not merely cultural; it is written into the guidance that governs the field, and reading that guidance precisely matters.

The United States federal guidance for operational technology, NIST SP 800-82 in its third revision of 2023, is explicit that active scanning is a hazard on a live network. It advises that owners exercise extreme caution when permitting active scanning on an operational network, on account of device sensitivity, and it recommends passive network monitoring, which captures and analyses traffic without injecting packets, as the baseline for visibility; active scanning is to be avoided unless thoroughly tested in the specific environment [11]. The guidance frames this as a deliberate tailoring of IT-normal controls to an OT context in which safety and continuity weigh as heavily as confidentiality [11]. IEC 62443 sits alongside it, deriving control requirements from zones and threat models rather than from a generic scanning practice imported from IT [3].

The guidance is not static, and the direction of its evolution argues against reading the standards as a blanket prohibition on active methods. Commentary in the current revision cycle has urged the standards bodies to move past a binary scan-or-do-not-scan framing, to establish passive assessment as the safe baseline and then define a formal risk-assessment framework under which controlled active scanning can proceed in graduated phases [12]. This is the same conclusion the technical picture points to: passive as the default, active as a deliberate, risk-assessed addition rather than a routine practice or a forbidden one.

The standards, in other words, do not resolve the choice by banning one method. They establish a default and a burden of proof. Passive monitoring is the baseline a defender is expected to have; active methods carry the burden of demonstrating that they are safe for the specific devices before they are used.

6. The encryption dimension

There is a further axis of difference that the traditional discussion predates, and on which primary measurement, rather than general argument, is now available for at least one modern encrypted control channel [13].

The naive framing is that a passive observer loses the payload under encryption while an active method, able to authenticate and participate in the session, does not. This is true as far as it goes, but hardware measurement on a Siemens S7-1200 G2 Web API carried over TLS 1.3 sharpens it in two ways that matter for OT specifically.

First, the passive blindness under strong encryption is close to total, not partial. On that channel a strictly passive observer on a mirror port cannot recover the method, the target of an action, the value written, or the process variable whose trajectory continuous monitoring would depend on; under TLS 1.3 even the server certificate is encrypted, so the endpoint's identity is not passively recoverable either [13]. What survives is metadata alone: the size, direction, and timing of otherwise opaque records.

Second, and less obviously, the sensor that does survive encryption fails against precisely the case that matters most. A metadata sensor keyed on record volume and timing detects bursts of activity, but an isolated, well-chosen write is nearly invisible to it: in a controlled test, three of six isolated writes produced no detectable excursion above the polling baseline, with no discernible pattern to which were seen, because an individual write adds only on the order of a dozen bytes to a continuous polling stream [13]. The activity a metadata sensor detects well is not the activity a targeted adversary generates. Under strong encryption, the class of passive signal that remains is the weakest against the single deliberate action.

The correct way to state the axis is therefore not that encryption favours active over passive, but that it moves the vantage. The evidence a continuous monitor depends on is available on a clear channel from a passive mirror; on an encrypted channel it is available only where the transport is legitimately terminated, at a gateway, a proxy, or a co-located component at a point that already decrypts [13]. Whether that relocation is acceptable is a question each deployment answers against its own constraints, and it applies to any observing method that depends on reading the traffic, passive or active alike.

Two things bound this correctly. The encrypted share of the installed base is still the minority, as the companion landscape briefing on S7 sets out; across most deployed industrial traffic today the payload remains readable [14]. But the direction of travel is towards more encryption, and each increment moves the balance on this axis over equipment that will still be in service long after the choice is made. Encryption is not a reason to prefer one method over another; it is a dimension on which the location of the observation, more than the method of it, is what must change.

7. The limits of any single method

The picture that assembles from the preceding sections is not a verdict for one method over the other. It is a set of trade-offs that fall differently along each axis, and a field in which the honest position is that no single method is complete on its own.

The industry has largely converged on this in practice. The mature pattern is hybrid: passive monitoring as the safe baseline a defender is expected to maintain, with controlled, OT-safe active queries added deliberately where the passive picture has gaps, scheduled and risk-assessed rather than run as routine [5] [12]. The standards point the same way, establishing passive as the default and placing the burden of proof on active methods to show they are safe for the specific devices before use [11]. Neither the technical picture nor the guidance supports a purist position in either direction.

Yet convergence on hybrid does not dissolve the underlying question; it reframes it. Each method has a defined horizon. The passive method sees what talks, and reads what the wire carries in clear; its horizon is drawn by what devices transmit and by whether that transmission is legible. The active method reaches

what is silent and recovers what must be asked for; its horizon is drawn by what the devices can safely be made to answer. Encryption redraws both, and moves the useful vantage from the mirror to the point of legitimate termination for anything that depends on reading the payload.

The open question that remains is whether a given method's horizon is sufficient for what a given environment asks of it. That is not a question a method answers in advance; it is one each deployment answers for itself, against the devices it actually runs, the traffic they actually generate, the encryption they are actually migrating towards, and the adversary it is actually concerned with. A method sufficient for an estate of clear, chatty, legacy controllers may fall short against one migrating to encrypted channels and worried about the single deliberate action rather than the noisy intrusion. The value of stating the trade-offs precisely is not that they resolve the choice, but that they let each defender see where the horizon of their chosen method lies, and judge honestly whether it reaches far enough.

References

[1] E. Samanis, J. Gardiner, and A. Rashid. A Taxonomy for Contrasting Industrial Control Systems Asset Discovery Tools. arXiv:2202.01604, 2022 (also published as SoK, ARES 2022).

[2] Field reports on OT devices that fault under IT scanning tools. PLCTalk practitioner forum, 2023.

[3] International Electrotechnical Commission. IEC 62443, Security for Industrial Automation and Control Systems (series).

[4] IOActive. Rethinking the CIA Triad in Operational Technology Environments. IOActive, 2026.

[5] Dragos. Passive Monitoring and Active Collection for a Complete OT Asset Inventory. Dragos, 2025.

[6] Nozomi Networks. How to Ensure You Have Complete OT/ICS Asset Visibility. Nozomi Networks, 2026.

[7] Tenable. How to Tackle OT Challenges: Asset Inventory and Vulnerability Assessment. Tenable, 2023.

[8] runZero. Active Scanning of Industrial Control Systems Safely. runZero, 2024; and Radiflow, on tap/span coverage gaps and baseline accumulation, 2023.

[9] runZero. Rules for safe active scanning in ICS environments (RFC-compliant traffic, no security probes, per-host rate limiting). runZero, 2024.

[10] FireMon. Active vs. Passive Scanning: Key Differences. FireMon, 2025.

[11] National Institute of Standards and Technology. SP 800-82 Revision 3, Guide to Operational Technology (OT) Security. NIST, September 2023.

[12] Reporting on NIST SP 800-82 revision commentary urging a graduated, risk-assessed approach to active scanning with passive as the safe baseline. OT.today, 2026.

[13] B. Salmazo. Contextual Evaluation under Encryption. Liscere Technical Report LTR-2026-04, 2026. https://liscere.com/assets/papers/liscere-evaluation-encryption.pdf

[14] B. Salmazo. Security by Default, Legacy by Deployment: The Encryption Gap in the Installed S7 Base. Liscere Briefing LB-2026-04, 2026.