


2026 FIELD GUIDE

AI-Powered OT Penetration Testing

What's Changing for Manufacturers



2026 FIELD GUIDE

Table of Contents

How artificial intelligence is reshaping operational-technology penetration testing — what changes, what stays human, and how manufacturers stay ahead.

01 — THE SHIFT TOWARD AI-ASSISTED SECURITY TESTING

- Why Manufacturing Became a Primary Target
- The Limits of Manual OT Penetration Testing
- Where Artificial Intelligence Enters the Workflow
- What “AI-Powered” Actually Means in Testing

02 — HOW AI CHANGES THE PENETRATION TESTING LIFECYCLE

- Intelligent Asset Discovery and Passive Analysis
- Machine-Speed Vulnerability Correlation
- Automated Attack-Path Modelling
- AI-Assisted Exploitation Under Guardrails
- From Findings to Prioritised Risk Narratives

03 — THE MANUFACTURER’S ACCELERATED THREAT LANDSCAPE

- Adversaries Are Adopting AI Too
- AI-Generated Malware and Adaptive Attacks
- The Expanding Attack Surface of the Smart Factory
- Deepfakes, Social Engineering, and Access

04 — GUARDRAILS: SAFETY, TRUST, AND CONTROL OF AI

- Keeping AI Testing Safe for Live Production
- False Positives and Hallucinations
- Human-in-the-Loop and Accountability
- Data Governance and Model Trust
- Where AI Must Not Operate

05 — BUILDING AN AI-READY OT SECURITY PROGRAM

- Selecting AI-Augmented Testing Capabilities
- Integrating AI Findings into Defensive Action
- Measuring What AI Testing Improves
- The Manufacturer’s AI-OT Readiness Model

REFERENCE

- AI, OT & Security Glossary
- References, Authors, and Contributors
- About DivIHN

Human-led. AI-accelerated. testing built for the modern manufacturer

EXECUTIVE PERSPECTIVE


The Tester Now Has a Machine Partner

For manufacturers, the operational network has become both the most valuable and the most exposed part of the business. Attacks are faster and more frequent, the smart factory keeps widening the attack surface, and skilled OT security testers remain scarce. Manual penetration testing — slow, periodic, and hard to scale — is struggling to keep pace.

Artificial intelligence is changing that equation. Used well, it accelerates the parts of penetration testing that machines do best — discovery, correlation, attack-path modelling, reporting — while leaving judgement, safety, and accountability firmly with the human expert. This is augmentation, not automation of the hacker.

AI does not replace the OT penetration tester. It gives one expert the reach of a team — and gives the manufacturer testing that can finally keep up with change.

THE NEW TESTING PARTNERSHIP



#1

MOST-ATTACKED INDUSTRY

Manufacturing has led all sectors for cyberattacks for four consecutive years.

<1 in 5

OT ROLES EASILY FILLED

Skilled OT security testers remain in short supply as demand climbs.

24/7

CHANGE NEVER STOPS

Continuous digitisation makes point-in-time testing obsolete faster.

THE SHIFT TOWARD AI-ASSISTED SECURITY TESTING

Why Manufacturing Became a Primary Target

Manufacturers now sit at the top of nearly every attacker's target list. The reason is simple economics: a stopped line costs money by the minute, which makes production the most effective form of leverage an adversary can hold. That pressure, combined with rapid digitisation and decades of legacy equipment, has turned the factory floor into a high-value, high-exposure target.

25%+

OF ALL CYBERATTACKS

Manufacturing absorbs the largest share of incidents of any industry.

\$260K+

PER HOUR OF DOWNTIME

Typical cost of unplanned production stoppage — the attacker's leverage.

70%

OF OT ATTACKS VIA IT

Most operational impact still begins on the enterprise side of the network.

WHY THE FACTORY FLOOR IS SO ATTRACTIVE



Downtime intolerance

Production cannot pause, so ransom pressure is at its highest here.



Converged IT and OT

Business and plant networks now connect, opening IT-to-OT paths.



Legacy equipment

Controllers built without security run for decades on the line.



The smart factory

IIoT, robotics and cloud MES add new, less-understood exposure.



Valuable IP

Designs, recipes and process data are prime targets for theft.



Supply-chain role

One manufacturer's outage ripples across many downstream customers.

THE SHIFT TOWARD AI-ASSISTED SECURITY TESTING

The Limits of Manual OT Penetration Testing

Manual penetration testing remains essential, but on its own it can no longer cover a modern manufacturing environment. It is slow, periodic, and constrained by a shortage of engineers who understand both offensive security and the safety realities of a live plant. The result is coverage that lags behind how fast the environment changes.

WHERE MANUAL TESTING STRUGGLES

Constraint	What it means on the plant floor
Point-in-time	A test is a snapshot; the environment has changed by the time the report lands.
Slow to scale	Large, multi-site operations exceed what a small expert team can cover by hand.
Skills scarcity	Few testers combine offensive skill with OT and safety knowledge.
Coverage gaps	Time limits force sampling; assets and paths go unexamined.
Safety caution	Fragile systems and live processes limit how deep manual testing can safely go.

3.5M

UNFILLED CYBERSECURITY ROLES WORLDWIDE
The talent gap is widest exactly where OT and offensive expertise overlap — the skills an industrial assessment needs most.

1–2×

TESTS PER YEAR, TYPICALLY
Most plants are assessed once or twice annually — while assets, vendors and connections change continuously.

THE SHIFT TOWARD AI-ASSISTED SECURITY TESTING

Where Artificial Intelligence Enters the Workflow

AI does not bolt onto testing as a single feature; it accelerates specific phases of an established process. The scoping, safety decisions and final judgement stay human. What changes is the speed and breadth of the analytical work in between — the discovery, correlation and modelling that used to consume most of a tester's time.

THE LIFECYCLE — AND WHERE AI ACCELERATES IT

1 Scoping & guardrails

Human-led — consequence, boundaries, rules of engagement.

2 Asset discovery · AI

Passive classification of assets, protocols and communications.

3 Vulnerability correlation · AI

Machine-speed matching against CVEs and threat intel.

4 Attack-path modelling · AI

Graph analysis of routes from entry to consequence.

5 Validation & exploitation

Human-approved, guardrailed, often in simulation.

6 Risk narrative & report · AI

Drafted by AI, reviewed and owned by the tester.

Human judgement bookends the process — setting scope at the start and owning the decision at the end. AI compresses the analytical middle from weeks to hours.

THE SHIFT TOWARD AI-ASSISTED SECURITY TESTING

What “AI-Powered” Actually Means in Testing

“AI-powered” is often used loosely. In practical OT penetration testing it refers to a handful of specific, mature techniques — not a single autonomous system. Understanding what each one does keeps expectations grounded and helps a manufacturer judge a claim honestly.

THE TECHNIQUES BEHIND THE TERM

Technique	What it does	In OT testing
Machine learning	Finds patterns and anomalies in large datasets	Classify assets; flag abnormal traffic
Large language models	Read, summarise and generate language	Recon synthesis; draft risk narratives
Graph analytics	Reasons over connected relationships	Model attack paths across the network
Automation & orchestration	Runs repeatable tasks at scale	Continuous, consistent evidence gathering

What it is

- A force-multiplier for human experts
- Faster discovery, correlation and reporting
- Continuous rather than once-a-year

What it is not


- An autonomous attacker on production
- A replacement for OT and safety expertise
- Infallible — its output must be verified

HOW AI CHANGES THE PENETRATION TESTING LIFECYCLE

Intelligent Asset Discovery and Passive Analysis

Every assessment begins with knowing what is there — and in manufacturing, a large share of assets is typically undocumented. AI excels at this first, safest phase: learning from passively captured traffic to identify devices, fingerprint protocols and map how systems communicate, without ever transmitting to fragile equipment.

FROM PASSIVE TRAFFIC TO A LIVING ASSET MAP



What AI infers from the wire

- Device type, vendor and role
- Protocols and firmware indicators
- Communication relationships & baselines
- Shadow and forgotten assets

30%+

OF OT ASSETS ARE TYPICALLY UNKNOWN BEFORE ASSESSMENT


AI-assisted passive discovery closes that gap quickly — and keeps the inventory current between tests.

HOW AI CHANGES THE PENETRATION TESTING LIFECYCLE

Machine-Speed Vulnerability Correlation

Once assets are known, the question becomes which of them are exposed — and which exposures actually matter here. Correlating device, firmware and configuration details against a constantly changing flood of advisories is exactly the kind of high-volume matching AI does far faster, and more consistently, than any human team.

FROM RAW SIGNALS TO WHAT MATTERS



AI narrows an unmanageable list to the exposures that are real, reachable and relevant to this plant.

MANUAL VS AI-ASSISTED CORRELATION


Dimension	Manual	AI-assisted
Speed	Days to weeks	Minutes to hours
Coverage	Sampled	Every known asset
Freshness	As of the last test	Continuous against new advisories
Consistency	Varies by analyst	Uniform, repeatable

HOW AI CHANGES THE PENETRATION TESTING LIFECYCLE


Automated Attack-Path Modelling

CA list of vulnerabilities is not a risk picture. What matters is whether an attacker can chain exposures into a route that reaches something critical. AI treats the environment as a graph and reasons over it – continuously mapping the paths from a plausible entry point to operational consequence, and keeping that map current as the network changes.

AN AI-MODELLED ATTACK PATH



MAPPED TO MITRE ATT&CK FOR ICS



```
graph LR; A[Initial Access] --> B[Discovery]; B --> C[Lateral Movement]; C --> D[Inhibit Response]; D --> E[Impair Process Control]; E --> F[Impact]
```

The diagram maps the attack path to MITRE ATT&CK for ICS, showing a sequence of tactics:


- Initial Access** → **Discovery** → **Lateral Movement** → **Inhibit Response**
- Impair Process Control** → **Impact**

HOW AI CHANGES THE PENETRATION TESTING LIFECYCLE

AI-Assisted Exploitation Under Guardrails

Proving that a path is real sometimes requires validation — and this is where AI must be most tightly constrained. On a live production line, exploitation is never handed to an autonomous system. AI can propose and help validate an approach, but action happens behind human approval, on agreed targets, and wherever possible in a lab or digital twin rather than on the plant itself.

THE APPROVAL GATE BEFORE ANY ACTION



AI may

- Suggest and rank validation options
- Run safely in a replica or twin
- Assemble and document evidence

AI may not


- Act autonomously on production
- Touch safety instrumented systems
- Proceed without human approval

HOW AI CHANGES THE PENETRATION TESTING LIFECYCLE

From Findings to Prioritised Risk Narratives

The last mile of testing has always been the slowest: turning raw findings into a clear, prioritised story a decision-maker can act on. Language models compress that work dramatically — drafting consequence-ranked narratives and remediation guidance in a consistent voice — while the tester remains the author of record, reviewing and correcting every claim.

THE REPORTING PIPELINE



Sample AI-drafted narrative

Critical · Internet-exposed remote access reaches the plant network with single-factor login, placing a controlled process one credential away from an outsider. Fix: broker access, enforce MFA, segment the jump host.

80%

LESS TIME ON REPORTING

Freeing expert hours for the judgement work that only a human can do — while every finding is still verified.

THE MANUFACTURER'S ACCELERATED THREAT LANDSCAPE

Adversaries Are Adopting AI Too

The strongest argument for AI-assisted testing is that attackers are already using AI to move faster. Reconnaissance, phishing, malware development and exploitation are all being accelerated. Defenders who test at manual speed are, in effect, bringing last year's pace to this year's adversary.

HOW ATTACKERS USE AI

- Automated reconnaissance at scale
- Convincing, tailored phishing in seconds
- Faster malware variants and evasion
- Rapid triage of stolen data and access
- Lower skill barrier for complex attacks

HOW TESTING MUST RESPOND

- **Match the tempo** with continuous, AI-assisted assessment
- **Model AI-enabled** attack paths, not just known ones
- **Test social engineering** including AI-generated lures
- **Reference MITRE ATLAS** for adversarial-AI techniques
- **Shorten the loop** from finding to fix

MITRE ATLAS extends the familiar ATT&CK model to attacks that target and abuse AI systems — a growing reference point as manufacturers add AI to both their operations and their defences.

THE MANUFACTURER'S ACCELERATED THREAT LANDSCAPE

AI-Generated Malware and Adaptive Attacks

AI lowers the cost of building and mutating malicious code. The immediate effect for manufacturers is speed and variety: more variants, generated faster, better at evading signature-based defences and at blending in with legitimate operational traffic. The malware need not be exotic to be effective — it only needs to arrive faster than defences adapt.

WHAT AI CHANGES ABOUT THE THREAT

**Polymorphism at scale**

New variants generated continuously to slip past signatures.

**Better evasion**

Code tuned to look like normal operational behaviour.

**Faster weaponisation**

Time from disclosed flaw to working exploit keeps shrinking.

**Living-off-the-land**

Legitimate tools abused, guided by adaptive automation.

**Tailored to the target**

Payloads shaped to a specific plant's stack and protocols.

**Lower skill barrier**

Capable attacks within reach of less-skilled adversaries.

Hours

NOT WEEKS — THE NEW EXPLOIT TIMELINE

As weaponisation accelerates, the window between disclosure and attack collapses — and periodic testing simply cannot see into it.

THE MANUFACTURER'S ACCELERATED THREAT LANDSCAPE

The Expanding Attack Surface of the Smart Factory

Every step toward the smart factory adds capability — and surface. Connected sensors, robotics, cloud production systems and digital twins create new pathways and dependencies, and the AI models running the factory become assets an attacker can target in their own right.

NEW SURFACE IN THE CONNECTED PLANT



IIoT & sensors

Thousands of small, long-lived devices at the edge.



Robotics & automation

Networked cells and controllers with physical reach.



Cloud MES & analytics

Production systems and data extending off-premises.



Digital twins

High-fidelity models that mirror — and expose — the process.



AI models themselves

Poisoning and manipulation of the models making decisions.



5G & private wireless

New connectivity that bypasses traditional boundaries.

When AI helps run the factory, the model becomes part of the attack surface. Model poisoning and manipulation of training or input data are exposures a modern assessment now has to consider.

THE MANUFACTURER'S ACCELERATED THREAT LANDSCAPE

Deepfakes, Social Engineering, and Access

Most intrusions still begin with a human being persuaded to grant access — and AI has made that persuasion far more convincing. Cloned voices and synthetic identities can now impersonate an engineer, a manager or a trusted vendor well enough to talk a help desk into a password reset or a remote-access session that reaches the plant.

A SYNTHETIC-IDENTITY PATH INTO OT



CONTROLS TO TEST

- Out-of-band verification for access and reset requests
- Vendor identity proofing beyond a voice or email
- MFA on every path into the operational zone
- Help-desk playbooks for suspected impersonation

WHY IT LANDS

- Trust in a familiar voice is hard to override
- Urgency short-circuits verification steps
- Vendors often hold standing remote access
- Synthetic media is now cheap and fast to produce

GUARDRAILS: SAFETY, TRUST, AND CONTROL OF AI

Keeping AI Testing Safe for Live Production

The arrival of AI does not change the first rule of OT testing: never compromise safety or availability to find a weakness. If anything, autonomy raises the stakes. AI in an industrial assessment is bounded by the same priorities that govern any human tester — and by explicit limits on what it may touch and how fast it may act.

GUARDRAILS ON AI IN PRODUCTION

- **Passive by default** — observe, don't transmit
- **Rate and scope limits** on any active step
- **Simulation first** — lab or digital twin
- **Human approval** before any action on production
- **Safety systems** excluded from AI-driven testing

THE PRIORITY THAT STILL GOVERNS

- 1 Safety**
No action may raise risk to people or process.
- 1 Availability**
Production continuity is protected throughout.
- 1 Integrity**
Control and data stay trustworthy.

Autonomy is earned in the lab, not on the line. On production, AI advises — a human acts.

GUARDRAILS: SAFETY, TRUST, AND CONTROL OF AI

False Positives and Hallucinations

AI is fast, not infallible. Machine-learning models raise false positives; language models can state something confidently that is simply wrong. In an office network a mistaken finding wastes time. In a plant, acting on one can stop a process – so every AI output is treated as a lead to verify, never a conclusion to act on.

EVERY AI OUTPUT PASSES A CONFIDENCE GATE



KNOWING THE FAILURE MODES


Failure mode	What happens	The control
False positive	A benign asset flagged as vulnerable	Evidence threshold before reporting
Hallucination	A confident but fabricated claim	Human verification of every finding
Stale context	Advice based on outdated data	Fresh, environment-specific inputs
Over-trust	Acting on output without checking	“Verify, then act” as policy

GUARDRAILS: SAFETY, TRUST, AND CONTROL OF AI

Human-in-the-Loop and Accountability

When AI is involved, the question of who is responsible becomes central. The answer does not change: a named human expert owns the scope, the decisions and the results. AI proposes; the tester disposes. Accountability cannot be delegated to a model, and a clear audit trail records who approved what, and why.

PROPOSE, REVIEW, DECIDE – THE LOOP



WHAT STAYS WITH THE HUMAN

- **Scope and rules** of engagement
- **Any action** on a live system
- **Final judgement** on risk and severity
- **Sign-off** and accountability for results

WHAT THE AUDIT TRAIL CAPTURES

- Which model produced which output
- The evidence behind each finding
- Who reviewed and approved it
- What action was taken, and when

GUARDRAILS: SAFETY, TRUST, AND CONTROL OF AI

Data Governance and Model Trust

Feeding operational detail into an AI system creates its own exposure. Network maps, asset inventories and findings are sensitive; where they are processed, who can see them, and whether they train someone else's model are questions a manufacturer must answer before adopting AI-assisted testing. The AI supply chain deserves the same scrutiny as any other.

CONTROLS FOR TRUSTWORTHY AI TESTING

**Data residency**

Know where operational data is stored and processed.

**On-prem option**

Keep sensitive analysis inside your own boundary where needed.

**No training on your data**

Contractual assurance your data won't train shared models.

**Model provenance**

Understand the model, its origin and its limits.

**Access & retention**

Least-privilege access and defined deletion timelines.

**Auditability**

A record of inputs, outputs and decisions.

QUESTIONS TO ASK AN AI TESTING PARTNER

- Where is our data processed and stored?
- Is our data ever used to train models?
- Can analysis run on-premises if required?
- How are model errors detected and handled?

GUARDRAILS: SAFETY, TRUST, AND CONTROL OF AI

Where AI Must Not Operate

Some boundaries are absolute. However capable the tooling becomes, there are places in an industrial environment where AI simply does not act alone — because the cost of a wrong move is measured in safety, not data. Naming these limits explicitly is part of doing the work responsibly.

THE HARD BOUNDARIES

**Safety instrumented systems**

Never an autonomous testing target — assessed by design review only.

**Control logic changes**

No AI-driven modification of how the process runs.

**Unsupervised active testing**

No autonomous exploitation on a live production line.

**Sole safety decisions**

A human always owns any judgement that affects safety.

The measure of a mature AI-OT programme is not what its tools can do — it is what they are deliberately not allowed to do.

BUILDING AN AI-READY OT SECURITY PROGRAM

Selecting AI-Augmented Testing Capabilities

Not every “AI-powered” claim is equal. For a manufacturer, the right capability is one that pairs genuine AI acceleration with deep OT expertise, strict guardrails and transparency about how it works. The technology matters less than the discipline around it.

WHAT TO LOOK FOR

Capability	Why it matters for a manufacturer
OT-native	Built for industrial protocols and safety realities, not repurposed IT tooling.
Human-led	Experts own scope, decisions and sign-off; AI augments them.
Guardrailed	Passive-first, approval gates, safety systems out of scope.
Transparent	You can see how a finding was reached and verify it.
Data-governed	Clear residency, retention and no-training commitments.
Continuous	Assessment that keeps pace between formal engagements.

A simple test of any claim: ask the provider to show how a specific finding was produced and validated.

Transparency about method is the clearest signal of a trustworthy capability.

BUILDING AN AI-READY OT SECURITY PROGRAM

Integrating AI Findings into Defensive Action

Faster findings are only valuable if they turn into faster defence. The speed AI adds to discovery should flow straight into a sequenced remediation plan – contain the reachable paths now, then invest in the structural controls that keep them closed – all mapped to recognised frameworks so progress is measurable.

0–30 days

- Close exposed remote access
- Enforce MFA on OT entry
- Isolate internet-facing services

30–90 days

- Segment critical zones
- Broker & log remote access
- Deploy OT monitoring

90–180+ days

- Continuous AI-assisted testing
- Mature detection & response
- Align to IEC 62443

FINDINGS MAP TO CONTROLS AND FRAMEWORKS

Finding theme	Defensive control	Framework anchor
Exposed remote access	Brokered access, MFA, just-in-time vendor accounts	IEC 62443 / NIST 800-82
Flat network	Zones & conduits, OT DMZ, segmentation	IEC 62443-3-3
No monitoring	Continuous OT detection and baselining	NIST CSF – Detect
AI/model exposure	Model governance, input validation	NIST AI RMF / MITRE ATLAS

BUILDING AN AI-READY OT SECURITY PROGRAM

Measuring What AI Testing Improves

Adopting AI-assisted testing is a decision that should be measured, not assumed. The value shows up in a handful of tangible metrics — broader coverage, faster validation, quicker remediation — and, critically, in the shift from a once-a-year snapshot to continuous assurance.

Full

ASSET COVERAGE
From a sampled subset toward the full known estate.

Hours

TO VALIDATE A FINDING
Down from days or weeks between discovery and action.

Live

CONTINUOUS ASSURANCE
Testing that tracks change rather than lagging behind it.

Measure	Manual baseline	AI-augmented
Asset coverage	Partial, sampled	Full known estate
Testing cadence	1–2× per year	Continuous
Time to validate	Days to weeks	Hours
Reporting effort	High, manual	Drafted, then reviewed
False-positive control	Analyst-dependent	Confidence-gated & verified

BUILDING AN AI-READY OT SECURITY PROGRAM

The Manufacturer's AI-OT Readiness Model

Readiness for AI-assisted OT security is a profile, not a single score. The scorecard below turns this guide into a self-assessment – where the programme is strong, where it is exposed, and where AI augmentation will add the most value first.

AI-OT READINESS SCORECARD

Domain	Maturity	The question to answer
Asset visibility	●●○○	Do we know every asset and connection into OT?
Continuous testing	●○○○	Are we assessed continuously, not once a year?
AI governance	●○○○	Do we govern how AI tools use our data and act?
Human oversight	●●○○	Does a named expert own every AI-assisted decision?
Data protection	●●○○	Is sensitive OT data kept in-boundary and untrained-on?
Detection & response	●○○○	Would we see – and handle – an OT intrusion?
Safety alignment	●●●○	Are AI limits and safety governed together?

Illustrative baseline – complete against your own environment. Four dots denote a mature, tested capability.

REFERENCE

AI & OT Security Glossary

A shared vocabulary for the AI and operational-technology terms used throughout this guide.

USING THE GLOSSARY

AI – Artificial Intelligence Systems that perform tasks normally requiring human intelligence, such as pattern recognition and language.	Adversarial ML Attacks that manipulate a model's inputs or training to change its behaviour.
ML – Machine Learning AI that learns patterns from data rather than being explicitly programmed.	Model poisoning Corrupting training or input data to make a model behave maliciously.
LLM – Large Language Model A model trained to read and generate natural language; used for recon synthesis and reporting.	Prompt injection Manipulating an LLM's instructions through crafted input.
Generative AI AI that produces new content – text, code or images – from learned patterns.	Digital twin A high-fidelity virtual replica of a process, useful for safe testing.
Agentic AI AI that can plan and take multi-step actions; in OT testing it is tightly guarded.	NIST AI RMF A framework for managing risks in the use of AI systems.
Anomaly detection ML that flags behaviour deviating from a learned normal baseline.	Deepfake Synthetic audio or video used to impersonate a real person.
Attack-path modelling Reasoning over a network graph to find routes from entry to consequence.	OT – Operational Technology Systems that monitor or control physical industrial processes.
Hallucination A confident but false output from a generative model; must be verified.	IIoT Industrial Internet of Things – connected sensors and devices on the plant floor.
False positive A benign item incorrectly flagged as a problem.	MES Manufacturing Execution System – software that runs and tracks production.
Human-in-the-loop A person reviews and approves AI output before it is acted on.	IEC 62443 The international standard for industrial control system security.
Guardrails Explicit limits on what an AI system may do and how far it may act.	Purdue Model A reference architecture placing systems in levels from field to enterprise.
MITRE ATLAS A knowledge base of adversarial techniques that target AI systems.	Continuous testing Ongoing assessment that keeps pace with a changing environment.

REFERENCE

References, Authors & Contributors

SELECTED SOURCES & FURTHER READING

- **MITRE ATLAS** — Adversarial Threat Landscape for AI Systems
- **NIST AI RMF** — Artificial Intelligence Risk Management Framework
- **IBM X-Force** — Threat Intelligence Index (manufacturing most-attacked)
- **CISA** — ICS advisories and OT security guidance
- **Dragos** — OT/ICS Cybersecurity Year in Review
- **IEC 62443 & NIST SP 800-82** — industrial security standards

METHODOLOGY NOTE

Figures in this guide are drawn from the public reporting above and from recurring patterns across industrial security assessments. Single indicative values illustrate a trend rather than any one organisation, and AI capabilities are described at the level of technique, independent of any specific product.

AUTHORS & CONTRIBUTORS

DivIHN OT Security Practice

Research and technical review by DivIHN's operational-technology, offensive-security and AI specialists.

Editorial & Design

DivIHN Marketing & Content team.

DISCLAIMER

This guide is provided for general information and executive awareness. It does not constitute specific security advice. AI-assisted OT testing must be planned and authorised for each environment, with appropriate safety and data-governance controls. DivIHN accepts no liability for actions taken solely on the basis of this document.

© 2026 DivIHN · AI-Powered OT Penetration Testing — 2026 Field Guide, first edition. All trademarks are the property of their respective owners.

REFERENCE

About DivIHN

DivIHN helps manufacturers secure the systems that keep production running. Our operational-technology security practice combines IT, OT and AI expertise — assessing industrial environments the way a modern adversary would, at machine speed, while protecting the safety and availability the business depends on.

HOW WE HELP MANUFACTURERS

AI-augmented OT VAPT

Human-led testing accelerated by AI, with guardrails built in.

Monitoring & detection

Continuous visibility and detection for operational networks.

Architecture & segmentation

Zones, conduits and secure remote access aligned to IEC 62443.

AI governance for OT

Data protection, oversight and safe adoption of AI in testing.

Ready to test at machine speed — safely?

Talk to DivIHN about AI-augmented OT penetration testing scoped around the consequences your operation cannot accept.

- www.divihn.com
- OT Security Practice



AI-Powered OT Penetration Testing

WHAT'S CHANGING FOR MANUFACTURERS · 2026 FIELD GUIDE