Decoy-Assisted Detection Metrics for Manufacturing Operational Technology Networks

Daniel Ward*
* ORCID: https://orcid.org/0009-0002-4466-6739
Department of Computer Science, Southern New Hampshire University, United States

DOI: 10.29322/IJSRP.16.06.2026.p17419
https://dx.doi.org/10.29322/IJSRP.16.06.2026.p17419

Paper Received Date: 16th May 2026
Paper Acceptance Date: 18th May 2026
Paper Publication Date: 22nd June 2026

Abstract- Manufacturing operational technology (OT) networks require detection methods that improve adversary visibility without disrupting physical processes. Deception technologies can generate high-confidence alerts because legitimate users should not interact with decoy workstations, historian tags, honeytokens, or simulated industrial services. However, organizations often struggle to determine whether deception deployments are effective after pilot installation. This paper develops a Decoy-Assisted Detection Metrics Model for manufacturing OT networks. Using design science and document analysis, the paper maps deception telemetry to detection efficiency, alert quality, adversary-technique observability, response readiness, operational burden, and decoy fidelity. The result is a practitioner-oriented measurement artifact with a metrics catalog, segment-specific scoring logic, and an illustrative synthetic application for engineering workstation, historian, remote-access, and controller-adjacent zones. The contribution is a non-human-subjects evaluation model that enables OT security leaders to measure deception value before, during, and after implementation.

Index Terms- cyber deception, detection metrics, industrial control systems, manufacturing cybersecurity, operational technology

I. INTRODUCTION

Manufacturing operational technology (OT) networks support production lines, batch processes, packaging equipment, quality systems, historians, engineering workstations, and controller-adjacent communications. Security monitoring in these environments is difficult because many controls that are routine in enterprise information technology (IT) environments can introduce latency, instability, unapproved network traffic, or operational risk when applied to industrial systems. NIST SP 800-82 Revision 3 emphasizes that OT security must account for performance, reliability, and safety requirements while addressing threats and vulnerabilities across programmable devices and systems that monitor or control the physical environment [2].

Deception technology provides complementary detection capabilities in this setting. Instead of relying solely on signatures, vulnerability scans, or generalized network anomaly detection, deception deploys controlled false assets or artifacts in locations where legitimate operators and production systems should not interact with them. A decoy engineering workstation, honey credential, historian tag, fake project file, or simulated protocol service can become a high-confidence indicator when it is touched. Prior dissertation research on deception technology adoption in manufacturing and critical infrastructure found that adoption barriers included compatibility concerns, resource limits, inadequate professional knowledge, infrastructure constraints, and concern about performance impacts [1]. Those findings suggest that organizations need more than a deployment checklist; they need defensible ways to measure whether deception is producing useful security outcomes.

The measurement gap is especially important in manufacturing. A decoy may be technically successful even if it does not detect a scan or an unauthorized login attempt. However, it may still fail as a program control if the alert is not routed, triaged, documented, and used to improve response readiness. Conversely, a low-volume decoy may provide significant value if it validates an access-control assumption, exposes a credential misuse path, or confirms that a remote-access segment is being probed. Traditional measures such as total alert count or number of deployed decoys do not capture those differences.

This paper addresses the measurement problem. The purpose is to develop a Decoy-Assisted Detection Metrics Model that manufacturing organizations can use to evaluate deception technology across OT network segments. Three questions guide the work: (1) Which metrics best describe the detection value of decoy-assisted monitoring in manufacturing OT networks? (2) How can deception telemetry be mapped to response and governance outcomes? (3) How can organizations score deception effectiveness across engineering workstations, historians, remote access, and controller-adjacent segments?

II. BACKGROUND

A. OT monitoring constraints

Manufacturing OT environments are not simply enterprise networks with industrial devices attached. They include long equipment lifecycles, specialized engineering software, deterministic communications, vendor-supported appliances, and strict uptime expectations. Monitoring must therefore minimize process interaction while still providing actionable evidence. The NIST Cybersecurity Framework 2.0 provides high-level cybersecurity outcomes for governance, identification,

protection, detection, response, and recovery, but it does not prescribe how organizations must achieve those outcomes [3]. A deception metrics model helps translate a specific technical capability into outcomes that leadership and auditors can understand.

B. Deception telemetry and adversary behavior

MITRE ATT&CK for ICS organizes adversary behavior into tactics and techniques such as initial access, discovery, lateral movement, collection, command and control, inhibit response function, impair process control, and impact [5]. Deception telemetry is valuable when it reveals technique-level behaviors that are otherwise difficult to distinguish from routine OT activity. Examples include interacting with a fake project file, attempting to use planted credentials, querying the canary historian tags, scanning simulated controller services, or accessing a remote-access lure.

C. Segment-specific detection opportunities

Engineering workstation zones are appropriate for file, credential, and project-archive lures because attackers often search for logic, drawings, and configuration material after gaining access. Historians and data platforms are appropriate for canary tags, bogus query paths, and false reporting artifacts because they bridge plant-floor data and enterprise visibility. Remote-access and OT demilitarized zones are appropriate for decoy jump hosts, vendor folders, and honey credentials, as adversaries and misconfigured tools frequently interact with remote management infrastructure. Controller-adjacent segments require the greatest care because decoys must never interfere with live control authority; however, passive simulated services can still help identify scanning, enumeration, and unauthorized message attempts.

D. Control and workforce alignment

Metrics also need to fit program governance. CISA Cross-Sector Cybersecurity Performance Goals 2.0 frames cybersecurity practices as outcome-driven actions for critical infrastructure owners and operators [4]. IEC 62443-2-1:2024 specifies security program policy and procedure requirements for industrial automation and control system asset owners and recognizes that many legacy systems require compensating measures [6]. The NICE Framework provides a common language for cybersecurity work, tasks, knowledge, and skills [7]. A manufacturing deception metrics program should therefore measure not only alerts, but also response readiness, ownership, training, maintenance burden, and evidence quality.

III. MATERIALS AND METHODS

A. Research design

This paper uses a non-human-subjects design science and qualitative document analysis. No interviews, surveys, experiments with human participants, or live production-system interventions were conducted. The design goal was an artifact that manufacturing organizations can use to evaluate deception effectiveness without exposing live OT processes to unnecessary risk.

B. Source selection

The analysis used current public standards, framework documents, and threat-modeling resources relevant to OT security and cybersecurity workforce practice. These included NIST SP 800-82 Revision 3, NIST CSF 2.0, CISA Cross-Sector Cybersecurity Performance Goals 2.0, MITRE ATT&CK for ICS, IEC 62443-2-1:2024, and the NICE Framework. The author's prior dissertation was used as the source of the adoption problem because it identified the ICS/OT deception barriers that motivate the measurement model [1].

C. Model construction

The model was constructed in four steps. First, common manufacturing deception placements were grouped into engineering workstation, historian/data, remote-access/OT demilitarized zone, and controller-adjacent segments. Second, measurable outcomes were derived from the detection, response, governance, and workforce functions emphasized across the selected frameworks. Third, the outcomes were translated into metrics that can be calculated from decoy logs, security information and event management (SIEM) events, incident tickets, packet captures, change records, and tabletop-exercise evidence. Fourth, the metrics were organized into a five-level scoring rubric.

D. Calculation method

The model uses normalized scoring rather than a single universal threshold because plants differ in architecture, production tempo, remote-access practices, and staffing. For each metric category, the organization assigns a score from 1 to 5 and may apply a local weight. The weighted category score can be summarized as follows: the total score equals the sum of each category score multiplied by its local weight, divided by the sum of all weights. This simple structure avoids false precision while still supporting repeatable comparison across quarters and sites.

E. Evidence artifacts

A metric is counted only when an evidence artifact supports it. Acceptable artifacts include decoy-platform logs, SIEM alerts, incident tickets, packet captures, remote-access logs, asset inventory records, change-control approvals, exercise reports, and post-incident review notes. This requirement reduces the risk that the model becomes a subjective maturity assessment without operational proof.

F. Illustrative application

The results include a synthetic example for demonstration only. The numbers are not field data and should not be interpreted as operational results from a specific plant. The example demonstrates how a manufacturing security team could translate decoy interactions into an effectiveness score after a pilot period.

IV. RESULTS

A. Metrics catalog

The model defines six metric categories. Detection efficiency focuses on speed and volume, but speed alone is not sufficient. Alert quality assesses whether deception telemetry produces actionable indicators rather than unmanaged noise. Technique observability links decoy interactions to ATT&CK for ICS behaviors. Response readiness evaluates whether the alert can be triaged and escalated. Operational burden ensures that deception

does not create an unsustainable maintenance load. Decoy fidelity evaluates whether the decoys are plausible enough to remain useful under adversary probing.

Detection efficiency. This category measures mean time to decoy detection, first-touch alert time, and decoy interaction rate by segment. In a manufacturing pilot, speed should be interpreted alongside operational context. A rapid alert from an engineering workstation lure may warrant immediate escalation, while repeated low-risk scans against a lab-isolated protocol decoy may be useful primarily for trend analysis.

Alert quality. This category measures high-confidence alert ratio, false-positive explanation rate, and duplicate-alert suppression. Because legitimate users should not touch deception assets, high-confidence alerting is a major advantage. Nevertheless, misconfigured backups, asset discovery tools, or authorized vendor activity can still touch decoys, so each alert should include an explanation path and closure reason.

Technique observability. This category measures how many ATT&CK for ICS techniques are directly or indirectly observable through deception telemetry. Useful examples include the discovery of remote systems, the use of valid accounts, interaction with project files, remote service probing, program download attempts, and suspicious historian queries. The objective is not to cover every tactic, but to select decoys that reveal likely adversary paths in the local plant architecture.

Response readiness. This category measures playbook activation rate, triage enrichment completeness, and escalation-path validation. A deception alert should identify the decoy touched, the source, the expected owner, the related network segment, likely ATT&CK mapping, and the response procedure. Without this context, even a technically accurate alert may not improve operational response.

Operational burden. This category measures monitoring effort, storage load, tuning time, maintenance frequency, and change-control effort. A deception pilot that requires constant manual tuning may fail even if it detects interesting activity. The model, therefore, treats manageable upkeep as a core effectiveness measure rather than an administrative afterthought.

Decoy fidelity. This category measures the realism of asset naming, protocol behavior, process values, access paths, and documentation lures. Fidelity does not require unsafe interaction with live control systems. It requires sufficient plausibility that an adversary, malware, or unauthorized user cannot immediately distinguish the decoy from its surroundings.

B. Evidence requirements

Each metric category should be tied to an observable evidence source. Acceptable sources include decoy logs, SIEM alerts, incident tickets, packet captures, firewall logs, remote-access logs, passive sensor records, asset inventory entries, change approvals, risk exceptions, tabletop results, and post-incident review notes. A metric is not counted unless at least one artifact supports it. This requirement reduces the likelihood that a deception pilot will be judged solely on informal impressions.

C. Scoring rubric

Each category can be scored from 1 to 5. A score of 1 indicates an ad hoc condition in which decoys exist but are not mapped to segments, owners, ATT&CK techniques, or response procedures. A score of 2 indicates a pilot condition in which selected decoys generate alerts, but metrics are limited to event counts and manual review. A score of 3 indicates a managed condition in which monthly metrics are collected and alerts feed ticketing or SIEM workflows. A score of 4 indicates an integrated condition in which deception telemetry supports ATT&CK coverage, OT response exercises, and governance evidence. A score of 5 indicates an optimized condition in which metrics drive placement changes, fidelity improvements, training updates, and continuous control validation.

The model allows local weighting. For example, a site with limited security staffing may weigh alert quality, response readiness, and operational burden more heavily than raw event volume. A site with a recent remote-access incident may weigh remote-access lures and credential honeytokens more heavily than historian tags. This flexibility helps the model support decision-making rather than forcing every plant into a single maturity template.

D. Illustrative synthetic application

A hypothetical 90-day pilot demonstrates the model. The plant places decoys in four manufacturing OT segments and reviews SIEM events, ticket outcomes, and OT change records. Engineering workstation lures receive a score of 4 because fake project-file access and honey credentials generate rapid, actionable alerts. Historian/data-layer lures receive a score of 3 because canary tags improve visibility, but baseline tuning is still required. Remote-access/OT DMZ lures receive a score of 4 because a decoy jump host and vendor-folder lure validate escalation paths. Controller-adjacent decoys receive a score of 3 because simulated protocol services detect scanning, but fidelity review and change-control evidence remain incomplete. The program average is 3.5, suggesting that the pilot is ready for managed expansion after tuning and change-control review.

V. DISCUSSION

A. Practical interpretation

The model shifts deception evaluation from “Did the decoy alert?” to “Did the decoy improve decision quality?” This distinction matters in manufacturing. A decoy that produces many alerts but no triage context may overload the security operations center. Conversely, a low-volume decoy that reliably identifies unauthorized credential use or engineering workstation reconnaissance may be more valuable than a noisy signature rule.

B. Governance value

The metrics also provide governance evidence. Leadership can see whether deception supports detection coverage, response exercises, and compensating controls for legacy systems. OT engineering teams can review whether decoy placement introduces any operational burden. Security teams can determine whether telemetry maps to ATT&CK for ICS and whether playbooks are triggered. Workforce leaders can use NICE-aligned task and skill language to assign ownership for monitoring, tuning, incident analysis, and continuous improvement.

C. Use in multi-site manufacturing

The model can also support multi-site comparison. A corporate OT security team could ask each site to score the same six categories every quarter. A plant with low event volume but strong readiness to respond may need different support than one with many alerts and weak triage. In this way, the metric model becomes a planning tool for training, staffing, and architecture decisions rather than a simple dashboard of security events.

D. Safety and change-control considerations

Decoy deployment in manufacturing should be governed through the same change-control discipline used for other OT security changes. The model assumes that decoys are passive or isolated, that they do not issue commands to live equipment, and that plant engineering personnel approve any placement that touches control-network segments. A metric that cannot be collected safely should be deferred or collected in a lab environment.

E. Limitations

The model is a design artifact rather than a field validation. The synthetic application illustrates scoring logic but does not establish statistical effectiveness. Organizations should validate the model through lab testing, cyber-range exercises, staged pilots, and post-incident evidence review. The model also assumes that decoys are isolated from live control authority and that deployment follows approved change-control processes.

F. Future work

Future research can extend this model by generating public synthetic datasets for industrial deception telemetry, comparing deception-derived alerts with intrusion detection system alerts, and testing whether specific decoy placements reduce mean time to detection. Sector-specific adaptations for water, power, transportation, and healthcare facility OT would also help determine which metrics generalize beyond manufacturing.

VI. CONCLUSION

Manufacturing OT security programs need detection methods that are useful, measurable, and operationally safe. Deception technology can provide high-confidence evidence when adversaries or unauthorized users interact with assets that legitimate operators should not touch. Yet the value of deception depends on measurement. This paper developed a Decoy-Assisted Detection Metrics Model that organizes deception telemetry into detection efficiency, alert quality, technique observability, response readiness, operational burden, and fidelity. The resulting artifact provides organizations with a practical way to evaluate deception pilots, explain the security value, and plan for safer expansion across manufacturing OT networks. The model is intentionally non-human-subjects and can be applied using logs, tickets, architecture records, and exercise evidence rather than new surveys or live adversary activity. Its practical value is that it lets OT security teams begin with a small pilot, document evidence, and expand deception only when the metrics show that the deployment improves detection and response without adding unacceptable operational burden.

VII. ACKNOWLEDGMENT

No external funding was received for this paper. The author acknowledges that the manuscript extends prior doctoral research on the adoption of deception technology in ICS/OT environments.

VIII. REFERENCES

[1] D. Ward, “Enhancing security: A comprehensive study on deception technology integration in manufacturing and critical infrastructure,” Ph.D. dissertation, University of the Cumberlands, Williamsburg, KY, USA, 2025. Available: https://www.proquest.com/dissertations-theses/

[2] K. Stouffer et al., Guide to Operational Technology (OT) Security, NIST Special Publication 800-82 Rev. 3, National Institute of Standards and Technology, Sep. 2023. doi: 10.6028/NIST.SP.800-82r3. Available: https://doi.org/10.6028/NIST.SP.800-82r3

[3] C. Pascoe, S. Quinn, and K. Scarfone, The NIST Cybersecurity Framework (CSF) 2.0, NIST CSWP 29, National Institute of Standards and Technology, Feb. 2024. doi: 10.6028/NIST.CSWP.29. Available: https://doi.org/10.6028/NIST.CSWP.29

[4] Cybersecurity and Infrastructure Security Agency, Cross-Sector Cybersecurity Performance Goals 2.0 for Critical Infrastructure, Dec. 2025. Available: https://www.cisa.gov/cross-sector-cybersecurity-performance-goals

[5] MITRE, “MITRE ATT&CK for ICS Matrix,” ATT&CK v19.1, Apr. 2026. Available: https://attack.mitre.org/matrices/ics/

[6] International Electrotechnical Commission, IEC 62443-2-1:2024, Security for industrial automation and control systems - Part 2-1: Security program requirements for IACS asset owners, Aug. 2024. Available: https://webstore.iec.ch/en/publication/62883

[7] R. Petersen, D. Santos, K. Wetzel, M. Smith, and G. Witte, Workforce Framework for Cybersecurity (NICE Framework), NIST SP 800-181 Rev. 1, National Institute of Standards and Technology, Nov. 2020. doi: 10.6028/NIST.SP.800-181r1. Available: https://doi.org/10.6028/NIST.SP.800-181r1

[8] National Institute of Standards and Technology, “NICE Framework: Current Versions,” updated May 18, 2026. Available: https://www.nist.gov/itl/applied-cybersecurity/nice/nice-framework-resource-center/nice-framework-current-versions

[9] National Institute of Standards and Technology, Security and Privacy Controls for Information Systems and Organizations, NIST SP 800-53 Rev. 5, Sep. 2020. doi: 10.6028/NIST.SP.800-53r5. Available: https://doi.org/10.6028/NIST.SP.800-53r5

[10] Cybersecurity and Infrastructure Security Agency, “Known Exploited Vulnerabilities Catalog,” 2026. Available: https://www.cisa.gov/known-exploited-vulnerabilities-catalog

IX. AUTHORS

First Author - Daniel Ward, Ph.D., Southern New Hampshire University, Department of Computer Science. Email: d.ward@snhu.edu. ORCID: https://orcid.org/0009-0002-4466-6739.

Correspondence Author - Daniel Ward, d.ward@snhu.edu.