Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 2540–2550
https://doi.org/10.54941/ahfe1008479

Designing an Experimental Environment for Training Stakeholder Orchestration by the Non-Expert One-Person SIRT in Manufacturing SMEs

Tomoyuki Kato[1] and Isao Kato[2]

[1]Manufacturing and Innovation Laboratory, Nagoya Institute of Technology, Gokiso-cho, Showa-ku, Nagoya, Aichi, Japan
[2]Nagoya Institute of Technology, Gokiso-cho, Showa-ku, Nagoya, Aichi, Japan

ABSTRACT

A cyber incident that stops a small or medium-sized enterprise (SME) in a Japanese manufacturing supply chain can stop its customers' lines within days. In many such firms the response falls to one non-technical employee, who cannot act on contact rules and decision criteria under pressure even when these exist. We call this person the one-person SIRT. This position paper describes the socio-technical design of an experimental environment for drilling that person. The competency we target is orchestration, a program-management term for aligning autonomous programs by presenting and selecting the criteria under which each judges its options, not by command. In incident response it means aligning several stakeholders' conflicting criteria into one defensible course of action under time pressure, which the one-person SIRT must do alone. Since no colleague is available, the drill gathers three to five one-person SIRTs from different firms into one team, a scaffold in which reasoning normally done alone is spoken and compared. AI stakeholder personas stand in for the missing departments, each holding private facts derived from a functional model of the firm, an affective profile, and a bounded authority to act on instructions. A state manager fires the events that change the situation. Reflection combines automated extraction, an expert retrieval system, and a human facilitator. The primary outcome is transfer to the individual, measured with a solo post-test on a new scenario. We present the design, its rationale, and a narrowed verification plan.

Keywords: Cyber incident response training, Tabletop exercise, Manufacturing SMEs, Supply-chain security, Stakeholder orchestration, Naturalistic decision making

INTRODUCTION

Cyber incidents in Japanese manufacturing travel along the supply chain. In 2022 a ransomware infection at Kojima Industries, a first-tier parts supplier, led Toyota to suspend all 28 lines at its 14 domestic plants (Toyota Motor Corporation, 2022). Many SMEs on such chains have no information systems department. Information technology (IT), the operational technology (OT) of the shop floor, and security are all handled by one non-specialist employee, often in general affairs or accounting, alongside other duties.

Received September 25, 2026; Revised November 3, 2026; Accepted November 15, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

This is the one-person IT department (*hitori joshisu*). A 2024 survey of 4,191 Japanese SMEs by the Information-technology Promotion Agency (IPA; renamed Innovation Platform Agency, Japan in June 2026) found that about 70% had no organized security structure, fewer than 40% had prepared even partial emergency procedures, and about 70% of those attacked in the previous three years saw the damage spread to business partners (IPA, 2025).

When an incident occurs, this one person is the whole response team. We call them the one-person SIRT (Security Incident Response Team), SIRT rather than CSIRT because in a factory the incident reaches physical production. The one-person SIRT is our trainee, and their situation defines the problem. There is no colleague to consult, no department whose position can be reconciled, and no occasion to rehearse. Paradoxically, then, training a person who will respond alone calls for a team. When three or four such people decide together, reasoning each would otherwise do silently is spoken, observed, and compared. Such a team can only be formed across firms, and we treat it as a scaffold for the individual rather than as the unit whose performance we claim.

Normative frameworks state what such a person should do: NIST CSF 2.0 in its Respond function and its supply-chain risk governance (NIST, 2024), IEC 62443 in foundational requirement FR 6 on timely response to events (IEC, 2013), and Japan's Cyber/Physical Security Framework (CPSF) (METI, 2019). None says how a non-expert acquires the ability. Nor does preparing contact routes and decision criteria in advance close the gap, since acting on a written procedure under pressure is itself a skill, acquired by rehearsal and lost without it (Arthur et al., 1998).

This position paper presents the socio-technical design of an experimental environment for that rehearsal. It offers a working definition of orchestration and a way of observing it, a cross-firm team used as a scaffold with transfer as the primary outcome, personas derived from a functional model of the firm, reflection divided among three roles, and a verification plan narrowed to one proposition.

PREVIOUS STUDY AND PROBLEM STATEMENT

Prior Work on Incident Response Exercises

Cyber incident response has the features that define naturalistic decision making: time pressure, high stakes, ill-structured problems, and uncertain, changing information (Klein, 2008). In such settings novices may perceive individual cues yet fail to integrate them into an understanding of what is happening (Endsley, 1995a). Tabletop exercises (TTX) train the decision-making and communication that technical ranges cannot (Angafor, Yevseyeva and He, 2020), including a virtual exercise for non-technical participants (Angafor, Yevseyeva and Maglaras, 2023) and an IT/OT incident-response exercise with critical-infrastructure operators in Japan, in which information sharing and decision-making across departments were observed (Nakayama, Koshijima and Watanabe, 2025). These formats assume several departments in the room whose positions the participants reconcile, or that the scenario author has pre-aligned. LLM-based exercises generate scenarios (Yamin

et al., 2024) or field autonomous agents that play the responding team (Liu, 2025); they generate the situation or automate the coordinating role. None targets a trainee who has no one to coordinate with.

Problem Statement

Standard incident handling, as codified in NIST SP 800-61 (NIST, 2025), presupposes severity classes, escalation rules, a communication plan, and assigned roles agreed in advance, so that the criteria for preferring one action to another are given. Most Japanese SMEs have not prepared these (IPA, 2025), and even where they exist they do not guarantee the ability to act on them: a rule to report to the parent firm within 24 hours does not say whether to report before or after restarting the line. For the one-person SIRT the criteria must be constructed during the incident, from stakeholders who state them only when asked, and in conflict. Kato and Koshijima (2026) argue that program management can bridge technical handling and management decisions; this design targets the competency that bridging requires. The problem is thus not the absence of a procedure but the untrained act of constructing and aligning criteria under pressure.

The decision situation to be reproduced therefore has a definite shape. An OT incident in such a firm places five parties in direct tension: production, which wants the line back; management, which weighs continuity against cost and liability; the external IT vendor, bound by contract and business hours; public relations, which wants no statement before the facts; and the customer, who wants a report and a delivery guarantee. Each holds facts the others lack and judges the same options (restart, isolate, pay, disclose) by a different criterion. Technical analysis constrains the options but cannot rank them; ranking requires that these criteria be elicited, their conflicts recognized, and one course of action chosen before the next development closes the options. That act is what the environment must train.

Orchestration: Definition and Functions

In the Project and Program Management (P2M) literature, orchestration is the core process of a multi-program platform: a higher governing body, the super-program, aligns programs that operate autonomously toward the corporate common value, not by command but by presenting and selecting the evaluation criteria under which each program chooses its own strategy (Kato, Kaede and Koshijima, 2016), a notion close to the asset orchestration of dynamic capabilities (Teece, 2007). Its functions are (i) presenting evaluation criteria and thereby giving strategic direction; (ii) mediating between parties whose languages, goals, and time horizons differ, such as management and the technical field, and keeping them dynamically aligned when the environment changes; (iii) reconfiguring assets so that the organization renews itself; (iv) revising the vision and criteria through feedback from practice; and (v) creating synergy no party could produce alone. In incident response the one-person SIRT must perform (i) and (ii) alone and within hours; (iv) appears across drills, when reflection revises the criteria; (iii) and (v) lie outside this paper. We therefore define orchestration in incident response

as the selection and time-pressured alignment of conflicting stakeholders' evaluation criteria into one defensible course of action. In decision-analytic terms the criteria are constructed from the concerns of those affected rather than found (Keeney, 1992). This separates orchestration from its neighbors, each of which takes the criteria as given: coordination manages dependencies between activities whose goals are fixed (Malone and Crowston, 1994), stakeholder management balances given interests, crisis communication chooses a message under reputational criteria, and team situation awareness builds a shared model of what is happening (Endsley, 1995a). The definition is provisional; what we hold fixed is the level at which it operates, the criteria themselves rather than actions under given criteria. How the competency is observed is specified with the evaluation.

DESIGN OF THE EXPERIMENTAL ENVIRONMENT

Figure 1 shows the environment. The team works at one terminal and a whiteboard; one member, the operator, types the team's queries and instructions to the personas. Each element is described below with the reason for it.

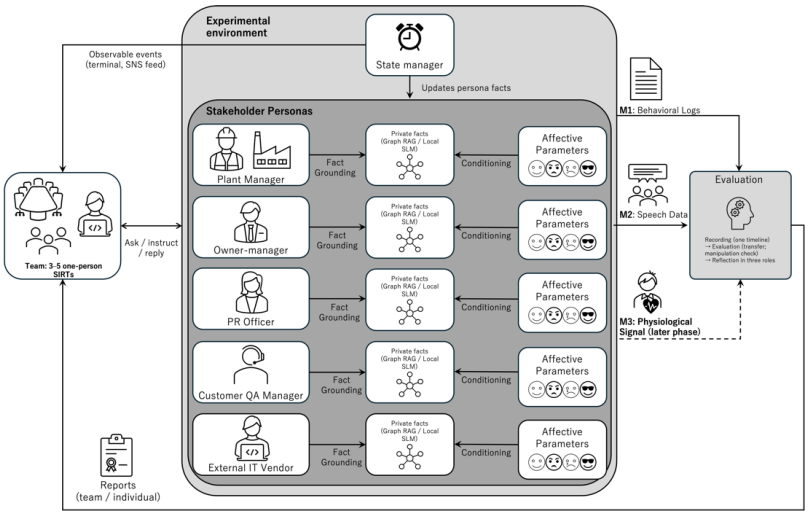


Figure 1: The experimental environment. No component reconciles stakeholder positions on the team's behalf.

The Team as a Scaffold, With a Solo Post-Test

Because such people are scarce in each SME, a team can be formed only across firms. The design premise, tested by the manipulation check of the evaluation section, is that when peers in the same situation must agree on whom to ask and what to prioritize, the judgments each normally makes alone are stated aloud and become comparable. The team is a scaffold (Wood, Bruner and Ross, 1976): a temporary support for what the learner cannot yet do alone, withdrawn as competence grows. Since the trainee will respond alone in reality, what must be shown is not that the team performed well but that the individual can perform without it; transfer, not team performance,

is therefore the primary outcome. It is measured by a written vignette given individually before and after the exercise (a different attack on the same firm), for which each participant writes a decision and its grounds, scored with the instruments of the evaluation section. The team has three to five members; four is the default, set by the number a facilitator can address individually in reflection. The operator issues each query or instruction only as the team's decided result, so that every logged act is a team decision.

Stakeholder Personas Derived From a Functional Model of the Firm

Five AI personas supply the departments a team of one lacks: the plant manager, the owner-manager, the person handling public relations, the customer's quality-assurance manager, and the external IT vendor; the in-house IT function is the trainee, not a persona. The vendor is included because about half of Japanese SMEs turn first to an external IT vendor when a security problem arises, and vendor-side shortfalls have contributed to incidents (IPA, 2026). A social-media feed of posts, rumors, leaks, and fabricated images, generated by the state manager described below, is an information source, not a persona.

Each persona holds private facts, an affective profile, and a bounded authority to act. In general, all three are derived from one functional model of the firm to be simulated: its functions, the role responsible for each, the contracts and reporting obligations that constrain each, and the information and material flows between them. A persona's private facts are the facts of the functions for which it is responsible; its authority is the set of those functions; an attack disables a function or an asset it depends on, and the loss propagates along the flows to dependent functions. Any notation that records responsibility and dependency suffices; the first instance is a parts machining firm of 50 to 100 employees, modeled with seven functions (order receipt, production planning, machining, inspection, shipping, invoicing, IT maintenance).

The knowledge is stored as a graph and retrieved with graph-based retrieval-augmented generation (Graph RAG), because what matters is relational (which contract binds which shipment, which function depends on which asset) and a graph store preserves it; replies are generated by a small language model running locally (a simulation of a real SME may contain its vulnerabilities), conditioned on the retrieved facts and on three parameters that set each persona's affective stance, rationality, anxiety, and assertiveness (Table 1); a persona's anxiety is raised when a later inject disables a function it is responsible for. The vendor's provisional contract (business-hours support for business systems only, next-business-day on-site response, emergency work billed separately) is an assumption: SMEs' low security spending, about 60% investing nothing (IPA, 2025), excludes round-the-clock service.

How the Scenario Evolves: Instructions and State Management

The scenario is driven by a state manager: not a role or a person but the component that changes the situation the team faces. It fires timed injects, exogenous events such as a second ransom note arriving or a fabricated image

of the attack appearing on social media, and it applies the consequences of the team's executed instructions. Each change is made in the functional model of the firm and propagates in two directions: to what the team can observe on its terminal and in the social-media feed, and to the private facts of every persona whose functions are affected. The state manager thus acts on the world the team is in, never as a party within it.

Two further decisions determine how the scenario evolves and who drives it: what the team can change, and what the personas may do on their own. Both follow from the requirement that the reconciliation be performed by the trainees and nowhere else. First, besides asking, the team may instruct a persona. Each persona's authority is the set of functions it is responsible for. An instruction within that authority is executed (the plant manager restarts the line, the vendor isolates the network, the owner-manager approves a payment or a report), and the scenario state changes accordingly: in the firm model the affected function or asset is disabled or restored, and the state manager updates the private facts of every persona whose functions depend on it. An instruction outside the authority, or one the persona resists because of its affective state, is not executed, and the persona replies why. The state manager changes the scenario only in these two ways, by timed injects and by executed instructions; it does not address the team as an interlocutor, summarize the personas, reconcile their positions, or recommend actions. Second, personas neither converse with one another nor initiate contact, although both happen in reality: if they exchanged information, part of the alignment would be done off-stage, and if they called first, they would set the agenda, whereas deciding whom to ask is the real difficulty. Each persona is a separate process with only the team's text as input, so leakage between personas is excluded, and because authority is enforced by the state manager rather than the model, prompt injection cannot execute actions; attempts are logged.

Scenario, Periodic Drills, and Reference Responses

The first scenario is a direct attack on OT, the situation for which IEC 62443 FR 6 sets the response requirements. A session covers the initial response, the first day of the incident compressed into two to three hours: the PLC of the main line stops in the morning; a ransom note appears on the production planning server before noon; the vendor's support conditions, the customer's delivery deadline, and a media inquiry reach the team in turn, as events observed rather than calls received; and by the end of the day the team must decide whether to restart (Table 1). Each inject is an exogenous event fired by the state manager; whom to query about it is the team's decision. Because trained skills decay with nonuse (Arthur et al., 1998), the environment is meant for periodic drills, and repetition requires scenario variants, which the functional model supplies by disabling different functions or assets of the same firm. Each inject carries an expert-authored reference response which forms the evidence model of the evaluation in the sense of evidence-centered design (Mislevy, Steinberg and Almond, 2003).

Table 1: Injects of the initial-response session, personas, and the criteria placed in tension.

Inject (Day 0)	Persona; Affective Setting	Criteria in Tension	CSF 2.0 / IEC 62443
Line halt (09:00)	Plant manager; rationality high, anxiety low	Manual restart vs. contain and preserve evidence	RS.AN-03; FR 6
Ransom note on planning server (11:00)	Owner-manager; anxiety high, rationality moderate	Continuity vs. cost and legality of payment; when to report to the parent firm	RS.MA-03, RS.MA-04; GV.SC-08
Vendor ticket auto-reply (13:00)	External IT vendor; rationality high, assertiveness low	Isolation takes hours, is billed, excludes controllers vs. restart now	RS.MA-02; RS.AN-03
Order system: delivery due tomorrow (15:00)	Customer QA manager; rationality high, assertiveness high	Report and delivery guarantee vs. the plant's recovery pace	GV.SC-05, GV.SC-08
Fake image on social media; media inquiry (16:30)	PR officer; anxiety high, assertiveness low	Disclose now vs. wait for facts	RS.CO-02, RS.CO-03
End-of-day decision (17:30)	All personas	Restart tomorrow vs. wait for the vendor's assurance	RC.RP-01

EVALUATION, REFLECTION, AND A NARROWED VERIFICATION PLAN

Three data streams are recorded on one timeline: behavioral logs (M1) of every query, instruction, reply, and state change; discussion transcripts (M2); and, in a later phase, physiological signals (M3) of stress that participants do not verbalize.

Observing Orchestration With Established Instruments

Orchestration is observed with three established instruments rather than a scheme of our own. First, the discussion transcript (M2) is coded by the functional perspective on group decision making (Hirokawa, 1985), whose requisite functions are problem analysis, establishment of evaluation criteria, generation of alternatives, and evaluation of their consequences; problem analysis and evaluation of negative consequences are the strongest predictors of decision quality (Orlitzky and Hirokawa, 2001). Orchestration appears in this coding as criteria establishment attributed to particular stakeholders and as consequence evaluation that weighs those criteria against one another; its timing within each inject window is examined by sequence analysis (Bowers et al., 1998). Second, because talk can be plentiful without insight, each

member's awareness of the stakeholders' criteria is probed directly: at inject boundaries the exercise is frozen and members answer queries on what each stakeholder currently wants and why, in the manner of SAGAT (Endsley, 1995b), scored against the personas' actual criteria; coverage against the expert's list of criteria follows the objective-generation paradigm of Bond, Carlson and Keeney (2008). Third, the adequacy of each decision is rated by two experts on behaviorally anchored scales adapted from emergency-team measures such as TEAM (Cooper et al., 2010), against the reference responses; the expert retrieval system is validated against these ratings.

Narrowed Verification Plan

The plan is narrowed to one primary proposition: participation in the environment, compared with a static TTX in which stakeholder positions are pre-aligned or absent, improves the individual's orchestration on a new scenario, measured by the pre/post change on the solo vignette (arm-by-time interaction). Criteria elicitation and reconciliation in the team (functional-perspective codes and probe accuracy) is the manipulation check and a candidate mediator, examined only correlationally since team format is constant across arms. Table 2 gives the outcomes by tier.

The primary outcome is individual-level, about 80 participants per year across two arms. Transfer effects of training are typically small (Blume et al., 2010), and with 40 per arm a two-sample comparison at the 5% level has 80% power only for a standardized difference of about 0.6 (Cohen, 1988), so smaller effects require accumulating cohorts. Phase 1 builds one persona; Phase 2 runs one team to establish inter-rater reliability of the coding and ratings and the expert retrieval system's agreement with human raters; Phase 3 is the comparison; Phase 4 is a six-month follow-up.

Table 2: Outcomes by tier in the narrowed verification plan.

Tier	Outcome	Measure	Data	Phase
Primary	Transfer to the individual	Pre/post change on solo vignette (functional-perspective codes and expert rating), arm-by-time	Vignette	3
Manipulation check	Criteria elicited and reconciled in the team	Functional-perspective codes per inject window; freeze-probe accuracy on stakeholders' criteria; sequence analysis	M2, probes	3
Secondary	Gap to expert reference	Behaviorally anchored expert rating vs. reference responses; expert retrieval system validated against two raters	M1, M2	2 to 3
Secondary	Standards coverage	Checklist of Table 1 subcategories addressed	M1	3

Reflection in Three Roles

Reflection after the exercise divides the work among three roles. Automated extraction reports what happened. An expert retrieval system, grounded in the standards (NIST, 2025; IEC, 2013) and the reference responses, compares the logged queries, instructions, and transcript with what an experienced responder would have done at each inject, and reports what the team accomplished and where it diverged. A human facilitator then asks what the team was thinking and why a view was not adopted, the questioning of taken-for-granted criteria that double-loop learning requires (Argyris and Schön, 1978); whether a language model can recover such value premises reliably has not been established, so this role is reserved for the facilitator. The retrieval system's validity as a rater is unproven and it is used only after the fact.

LIMITATIONS AND OPEN QUESTIONS

Several limitations bound what this paper claims. The environment has not yet been implemented or run; the one-person configuration is inferred from survey evidence on the absence of structure and awaits a descriptive study; the reference firm, the affective profiles, and the vendor's contract are assumptions specific to Japanese manufacturing SMEs; the first-year sample detects medium-to-large effects only; the distinction between orchestration and neighboring concepts is argued, not demonstrated. Beyond these limits, we put three questions to the human factors and systems engineering community. (1) Are the functional perspective's categories, developed for groups pursuing a shared goal, adequate when the criteria belong to conflicting stakeholders outside the group, and do freeze probes disturb a two-hour session? (2) Does a pre/post gain on a written vignette, which only approximates solo response, predict behavior in a later real incident? (3) Can the expert retrieval system's agreement with human raters be established, and does drilling against a fixed reference teach the rubric rather than the judgment?

CONCLUSION

We have presented an environment for drilling the one-person SIRT in orchestration: eliciting stakeholders' conflicting criteria and aligning them into one defensible course of action before the next development closes the options. Personas grounded in a functional model of the firm supply the missing departments; the reconciliation is left to the humans; evaluation is narrowed to transfer to the individual. We ask the community what the design lacks.

ACKNOWLEDGMENT

We are deeply appreciative of the funding provided by the Industrial Cyber Security Center of Excellence (ICSCoE) of the Innovation Platform Agency, Japan (IPA).

REFERENCES

Angafor, G.N., Yevseyeva, I., He, Y. (2020). Game-based learning: A review of tabletop exercises for cybersecurity incident response training. Security and Privacy, Volume 3 No. 6, e126.

Angafor, G.N., Yevseyeva, I., Maglaras, L. (2023). Scenario-based incident response training: Lessons learnt from conducting an experiential learning virtual incident response tabletop exercise. Information & Computer Security, Volume 31 No. 4, pp. 404–426.

Argyris, C., Schön, D.A. (1978). Organizational Learning: A Theory of Action Perspective. Reading, MA: Addison-Wesley.

Arthur, W., Bennett, W., Stanush, P.L., McNelly, T.L. (1998). Factors that influence skill decay and retention: A quantitative review and analysis. Human Performance, Volume 11 No. 1, pp. 57–101.

Blume, B.D., Ford, J.K., Baldwin, T.T., Huang, J.L. (2010). Transfer of training: A meta-analytic review. Journal of Management, Volume 36 No. 4, pp. 1065–1105.

Bond, S.D., Carlson, K.A., Keeney, R.L. (2008). Generating objectives: Can decision makers articulate what they want? Management Science, Volume 54 No. 1, pp. 56–70.

Bowers, C.A., Jentsch, F., Salas, E., Braun, C.C. (1998). Analyzing communication sequences for team training needs assessment. Human Factors, Volume 40 No. 4, pp. 672–679.

Cohen, J. (1988). Statistical Power Analysis for the Behavioral Sciences, 2nd ed. Hillsdale, NJ: Lawrence Erlbaum.

Cooper, S., Cant, R., Porter, J., Sellick, K., Somers, G., Kinsman, L., Nestel, D. (2010). Rating medical emergency teamwork performance: Development of the Team Emergency Assessment Measure (TEAM). Resuscitation, Volume 81 No. 4, pp. 446–452.

Endsley, M.R. (1995a). Toward a theory of situation awareness in dynamic systems. Human Factors, Volume 37 No. 1, pp. 32–64.

Endsley, M.R. (1995b). Measurement of situation awareness in dynamic systems. Human Factors, Volume 37 No. 1, pp. 65–84.

Hirokawa, R.Y. (1985). Discussion procedures and decision-making performance: A test of a functional perspective. Human Communication Research, Volume 12 No. 2, pp. 203–224.

IEC (2013). IEC 62443-3-3:2013. Industrial communication networks – Network and system security – Part 3-3: System security requirements and security levels. Geneva: International Electrotechnical Commission.

IPA (2025). 2024 survey on information security measures in small and medium-sized enterprises: Report (in Japanese). Tokyo: Information-technology Promotion Agency, Japan. Website: https://www.ipa.go.jp/security/reports/sme/sme-survey2024.html

IPA (2026). Report on the FY2024 program for strengthening the security response capability of regional IT vendors (in Japanese). Tokyo: Information-technology Promotion Agency, Japan. Website: https://www.ipa.go.jp/security/reports/sme/it-vendor.html

Kato, I., Koshijima, I. (2026). Fundamental Study on a Framework to Bridge the Disruption Between Technology and Management -A Proactive Incident Response Framework Based on P2M- (in Japanese). Journal of the International Association of P2M, Volume 20 No. 2, pp. 22–41.

Kato, I., Kaede, M., Koshijima, I. (2016). Decision-Making Protocol for Orchestration under P2M Framework (in Japanese). Journal of the International Association of P2M, Volume 11 No. 1, pp. 231–248.

Keeney, R.L. (1992). Value-Focused Thinking: A Path to Creative Decisionmaking. Cambridge, MA: Harvard University Press.
Klein, G. (2008). Naturalistic decision making. Human Factors, Volume 50 No. 3, pp. 456–460.
Liu, Z. (2025). AutoBnB: Multi-agent incident response with large language models. In: Proceedings of the 2025 13th International Symposium on Digital Forensics and Security (ISDFS). Piscataway, NJ: IEEE, pp. 1–6.
Malone, T.W., Crowston, K. (1994). The interdisciplinary study of coordination. ACM Computing Surveys, Volume 26 No. 1, pp. 87–119.
METI (2019). The Cyber/Physical Security Framework (CPSF), Version 1.0. (In Japanese), Ministry of Economy, Trade and Industry, Japan.
Mislevy, R.J., Steinberg, L.S., Almond, R.G. (2003). On the structure of educational assessments. Measurement: Interdisciplinary Research and Perspectives, Volume 1, No. 1, pp. 3–62.
Nakayama, K., Koshijima, I., Watanabe, K. (2025). Empirical study of information sharing and decision-making in IT/OT incident response. In: Applied Human Factors and Ergonomics International, Volume 199, pp. 2334–2344.
NIST (2024). The NIST Cybersecurity Framework (CSF) 2.0. NIST CSWP 29. Gaithersburg, MD: National Institute of Standards and Technology.
NIST (2025). Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile. NIST SP 800-61r3. Gaithersburg, MD: National Institute of Standards and Technology.
Orlitzky, M., Hirokawa, R.Y. (2001). To err is human, to correct for it divine: A meta-analysis of research testing the functional theory of group decision-making effectiveness. Small Group Research, Volume 32 No. 3, pp. 313–341.
Teece, D.J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. Strategic Management Journal, Volume 28 No. 13, pp. 1319–1350.
Toyota Motor Corporation (2022). Suspension of operations at domestic plants on March 1 (28 February 2022). Website: https://global.toyota/en/newsroom/corporate/3696105 1.html (Accessed 15 Sep. 2026)
Wood, D., Bruner, J.S., Ross, G. (1976). The role of tutoring in problem solving. Journal of Child Psychology and Psychiatry, Volume 17 No. 2, pp. 89–100.
Yamin, M.M., Hashmi, E., Ullah, M., Katt, B. (2024). Applications of LLMs for generating cyber security exercise scenarios. IEEE Access, Volume 12, pp. 143806–143822.