Trusted AI Decision Intelligence Support for Nuclear Material Accountability and Control

Frank L. Greitzer
Cogility Software
fgreitzer@cogility.com

Abstract: Despite significant investment in technical safeguards, the nuclear industry faces challenges detecting human-centric threats to nuclear material control and accountability (MC&A), including information theft, sabotage, nuclear material diversion, and unintentional errors. The high safety, regulatory, and geopolitical demands of the nuclear domain require more mature, human-centric risk models that exceed compliance-focused controls. Traditional perimeter defenses and cybersecurity tools often fail to identify early behavioral risk indicators, leaving analysts in a reactive posture. Effective solutions require proactive strategies that provide timely, reliable, and explainable analytic assessments. This highlights the need for a behavioral-science-driven approach that identifies behavioral indicators of stress, complacency, poor security culture, or malicious intent when assessing anomalies in user behavior, access patterns, and material handling.

Keywords: insider threat; nuclear cybersecurity; human factors; behavioral analytics; SOFIT; Expert AI; trusted AI; OT/ICS

INTRODUCTION

Insider risk remains one of the most difficult problems in nuclear security because trusted personnel possess legitimate access, operational knowledge, and the ability to exploit organizational blind spots from within. In nuclear environments, the potential consequences of insider actions extend beyond information loss to include sabotage, material diversion, operational disruption, and erosion of public confidence. International guidance has long recognized that insiders present a distinct threat category requiring measures that complement physical protection, cybersecurity, and material accounting controls [1], [2].

At the same time, the broader energy sector illustrates the scale and cost of the challenge. Public reporting indicates that insider incidents in energy and critical infrastructure are costly and increasingly complex, while research on operational technology and industrial control system security continues to show that human factors, weak security culture, and socio-technical vulnerabilities play central roles in adverse events [3]–[6]. These observations are especially salient for nuclear organizations, where mixed workforces, contractor dependence, tightly coupled operations, and expanding digital connectivity increase the need for continuous, human-centered risk assessment. This paper proposes that trusted AI and decision intelligence can strengthen insider risk management by synthesizing cyber, physical, personnel, and contextual data into explainable assessments that help analysts detect risk earlier and respond more proportionately.

CASES ILLUSTRATING BEHAVIORAL RISK PATTERNS

Historical cases from nuclear and closely related settings underscore that insider risk is not limited to a single motive or actor type. Instead, incidents span sabotage, theft, attempted espionage, policy violations, and severe human error, with precursors that may include disgruntlement, interpersonal conflict, probing

of controls, financial pressure, or unusual interest in restricted systems. Table 1 summarizes several illustrative cases and highlights the behavioral implications relevant to insider risk management.

Table 1. Illustrative Real-World Insider Risk Cases in the Nuclear Security Domain

Case	Description	Behavioral Implications
U.S. Army's SL-1 reactor (1961)	At the U.S. Army's SL-1 reactor in Idaho, rapid withdrawal of the central control rod caused a prompt critical excursion that killed the three operators on duty. The official investigation concluded that the evidence was consistent with rapid and extensive manual motion of the central control rod but did not establish a motive or intentional sabotage finding.	This is, at minimum, a cautionary case about **human and organizational factors** illustrating the consequences of single-point human action in a high-hazard system, and the importance of personnel reliability, procedural discipline, and design features that reduce vulnerability to unsafe manual acts [7], [8]. While more speculative, investigative news accounts revealed serious **job performance** concerns and **interpersonal issues**, and possible contributing factors related to **mental health issues, extreme stress, or suicidal ideation**.[9]
GE Wilmington Fuel Facility (1979)	In 1979, a temporary contractor employee at the General Electric fuel-fabrication facility in Wilmington, North Carolina, removed low-enriched uranium and attempted to extort money for its return. NRC notices describe how he exploited access weaknesses, including use of a driver's license resembling a badge and a vulnerable access path, before the material was recovered and the individual arrested.	This case demonstrates "probe and test" behavior by a knowledgeable insider. It emphasizes that behavioral monitoring must look for **unusual interest in security** or repeated, **unexplained access** to secure areas outside of one's normal duties. [10], [11].
San Onofre Diesel Generator Sabotage (2012)	In 2012, Southern California Edison reported possible tampering involving coolant found in the oil system of an emergency backup diesel generator at the San Onofre Nuclear Generating Station. Public reporting and NRC correspondence indicate that sabotage was investigated, but the publicly available sources located for this paper do not establish a confirmed perpetrator or motive.	This case is useful as a suspected insider tampering scenario. Reporting and NRC communications indicated suspicion of a **disgruntled employee**, with low morale and labor tensions noted. [12], [13] It highlights the need to correlate equipment anomalies with maintenance access, workforce stressors, and contemporaneous organizational tensions.
NRC and DOE	Former DOE and NRC employee Charles H. Eccleston attempted to sell NRC employee email-address data to people he believed were acting for a foreign government and later sent spear-phishing emails targeting DOE personnel. Official case materials indicate the conduct was financially motivated and included efforts to facilitate cyber compromise of U.S. government systems associated with nuclear-related missions [14], [15].	This case illustrates insider-enabled **cyber** and **espionage** risk, including misuse of privileged knowledge, attempted **recruitment** by a foreign actor, and **data exploitation for financial gain**. It also reinforces the need for continuous evaluation and equivalent screening standards for contractors and former insiders with residual knowledge of systems and personnel [14], [15].
Doel-4 Belgium (2014)	At Belgium's Doel-4 plant in 2014, the reactor shut down automatically after a loss of turbine lubrication in the non-nuclear part of the station. The Belgian nuclear regulator reported that the evacuation valve for the lubricating oil had been opened manually without instruction, leading to a sabotage investigation; the perpetrator was not identified.	This case underscores the challenge of trusted insider access in critical areas and the value of layered mitigations such as additional cameras, stronger badge controls, and broader application of the "four-eyes" principle for sensitive zones [16], [17],[18].

Taken together, these cases suggest that effective nuclear insider risk programs should emphasize four recurring needs: (a) **Behavioral indicators** such as disgruntlement, unexplained access behavior, abrupt changes in performance, financial strain, or unusual external contacts may provide early warning when interpreted in context. (b) **Continuous evaluation** is more effective than periodic screening alone, particularly for contractors, temporary staff, and personnel with privileged access. (c) **Strong security culture** and psychologically safe reporting channels are important to ensure that employees understand expectations, trust reporting channels, and perceive intervention as protective rather than purely punitive. (d) **Attention to organizational weaknesses** that create opportunities for misuse, concealment, or escalation.

Technical controls and a compliance posture are necessary but insufficient unless paired with governance, whole-person behavioral analytics, cross-functional coordination, and organizational resilience.

BARRIERS TO A WHOLE-PERSON INSIDER RISK APPROACH

Although the case for whole-person insider risk management is strong, implementation remains difficult in practice. Many organizations still manage personnel reliability, cybersecurity, physical security, and incident response through separate systems and organizational units, preventing analysts from developing a coherent view of risk. This fragmentation is compounded by privacy concerns, uncertainty about data governance, limited analytic capacity, and the challenge of distinguishing benign stress or error from indicators of elevated concern. In high-consequence nuclear settings, these barriers are magnified because programs must satisfy rigorous legal, regulatory, and ethical expectations while maintaining workforce trust [1], [19], [20]. Key barriers include:

- **Fragmented data and organizational silos:** HR, security, cyber, legal, and operational data are often not integrated in a way that supports timely, contextualized assessment.
- **Governance and ownership gaps:** no single authority may be accountable for enterprise-wide people risk, escalation thresholds, and intervention protocols.
- **Privacy, trust, and acceptability concerns:** monitoring programs can be perceived as surveillance if purpose, safeguards, and limits are not clearly communicated.
- **Resource and tooling constraints:** many organizations lack mature analytic platforms capable of integrating behavioral, cyber, and operational context.
- **Reactive operating models:** programs often emphasize incident response after a policy violation or alert rather than early intervention based on patterns of concern.

These constraints help explain why many insider programs remain compliance-oriented and reactive. They also clarify the potential value of trusted AI: not as a replacement for expert judgment, but as a means to synthesize diverse evidence, surface interpretable patterns, and support earlier, more proportionate decisions.

ENHANCING NUCLEAR INSIDER RISK MANAGEMENT

Trusted AI can strengthen insider risk management in at least three ways. First, it can synthesize large volumes of heterogeneous data—such as badge events, system activity, access requests, incident reports, training records, and selected personnel signals—to establish contextual baselines for individuals, roles, and teams. Second, it can help analysts recognize meaningful combinations of weak signals that would otherwise remain dispersed across multiple systems. Third, if properly designed, it can present those assessments in a transparent and auditable form that supports review, challenge, and accountable action rather than opaque automation [19], [21], [22].

For nuclear applications, trustworthy AI is especially important because performance alone is insufficient. Analytic outputs must also be understandable, traceable, secure, and appropriate for safety- and security-

critical decision environments. Guidance on trustworthy AI emphasizes characteristics such as validity, reliability, transparency, explainability, privacy enhancement, and human oversight [21]. Recent guidance on AI in insider risk management similarly stresses that organizations should treat AI as an augmentation capability, supported by governance, model validation, bias monitoring, and clear escalation procedures rather than as a self-acting replacement for multidisciplinary review [19].

Accordingly, the most promising role for AI in nuclear insider risk management is not generic anomaly scoring alone, but decision intelligence: an architecture that combines structured data fusion, domain knowledge, behavioral science, and human-in-the-loop adjudication to produce timely and defensible assessments.

A DECISION INTELLIGENCE EXPERT AI MODELING APPROACH

A decision intelligence platform provides a useful architectural foundation for insider risk management because it is oriented toward decision support rather than isolated alerts. In the nuclear context, such a platform can integrate cyber telemetry, physical access data, personnel and training information, incident reports, case management records, and selected organizational context to produce a continuously updated view of risk. This supports analyst triage by relating anomalies to role expectations, operational tempo, prior history, and relevant behavioral indicators rather than treating each signal as an independent event.

This approach has been implemented in the Cogynt.ai decision intelligence platform, which uses an expert AI framework based on hierarchical complex event processing (HCEP). Rather than relying exclusively on opaque machine learning outputs, the platform encodes analyst tradecraft, domain knowledge, and behavioral science concepts into layered event patterns that can be inspected and audited. This pattern-based approach is intended to support traceable inference, analyst trust, and more portable reasoning across scenarios than a purely statistical model trained on a narrow historical dataset [23].

Cogynt.ai models the reasoning process of experienced analysts by representing observables, intermediate assessments, and higher-level risk judgments as linked patterns. The objective is not to automate final decisions, but to help analysts detect concerning trajectories earlier, understand why a case was surfaced, and review the evidence trail supporting a notification. Because each assessment is grounded in explicit pattern logic and supporting evidence, the platform is designed to provide explainability by construction, thereby reducing the "black box" problem that often limits operational trust in AI-enabled systems.

The platform incorporates a whole-person assessment perspective using diverse structured and unstructured inputs that may include user activity monitoring, access-control events, incident reports, personnel data, and relevant external indicators (e.g., publicly accessible financial data, law enforcement records, social media) when authorized. HCEP then evaluates these observables against a hierarchy of expert-defined patterns. Lower-level observations are mapped to potential risk indicators (PRIs), which can combine into more abstract behavioral or operational patterns and, ultimately, into threat-relevant risk profiles such as sabotage, fraud, espionage, etc. The result is a dynamic behavioral risk picture that can support continuous triage and case management.

Figure 1 shows the Cogynt.ai workflow that begins with continuous ingestion of heterogeneous data streams. These inputs are evaluated for candidate indicators and then aggregated into multi-stage patterns representing combinations of psychosocial, cyber, physical, and organizational signals. When the resulting assessment crosses an analyst-defined threshold, the system generates a notification with provenance that allows the analyst to inspect the evidence, create a case record, collaborate with other subject matter experts, and determine whether mitigation or additional review is warranted. In this sense, the platform is designed to "shrink the haystack" by prioritizing interpretable cases rather than flooding analysts with disconnected alerts.

Figure 1. Cogynt.ai Process Flow

A central feature of this approach is analyst oversight. Experts can review, validate, and refine the patterns used by the model, examine the specific events that contributed to an assessment, and incorporate organization-specific priorities or constraints. This creates a feedback loop in which the analytic model remains auditable and adaptable while preserving human accountability for consequential decisions.

Figure 2 shows a conceptual diagram of the Cogynt.ai HCEP behavioral analytic process. The ingest layer of the Cogynt.ai platform receives structured and unstructured data feeds, including technical host/network monitoring and behavioral data, which comprise the "observables" that are available for all individuals (entities) within the enterprise. The first stage of analysis concerns identification of PRIs that are associated with collections of observables. For example, a certain combination of host/network data can be recognized as a print job sent to a network printer; an incident report may contain one or more observables that are associated with a Human Resources event such as a promotion or various types of disciplinary actions.

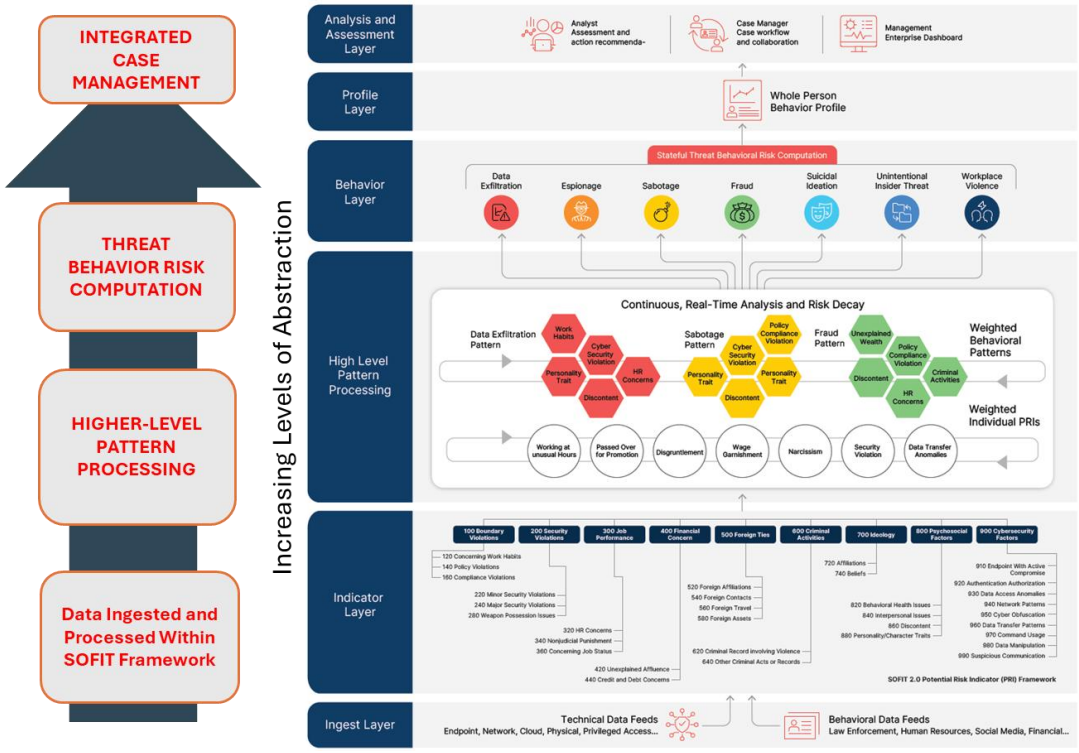

Figure 2. Cogynt.ai Hierarchical Complex Event Processing Behavioral Analytic

A key resource is an underlying knowledge base of PRIs. In the figure, the Indicator Layer shows the higher levels of the SOFIT (Sociotechnical and Organizational Factors for Insider Threat) knowledge base of insider risks [24], which provides a useful taxonomy of several hundred PRIs, organized into nine classes as shown in the diagram. At the lowest level of the taxonomy (not shown) are the PRIs, which fall into various subclasses. For example, the SOFIT class, Boundary Violations, comprises three subclasses (Concerning Work Habits, Policy Violations, and Compliance Violations); the Concerning Work Habits subclass (see Figure 3) contains seven PRIs: Working at Unusual Hours, Poor Time Management, Blurred Professional Boundaries, Nonproductive Socialization, Lack of Confidentiality, Reconnaissance, and Cyberloafing). Each of the SOFIT PRIs is in turn linked (with varying "weights" or strengths of association) to different insider risk behaviors of concern, shown in the Behavior Layer of the HCEP diagram in Figure 2.

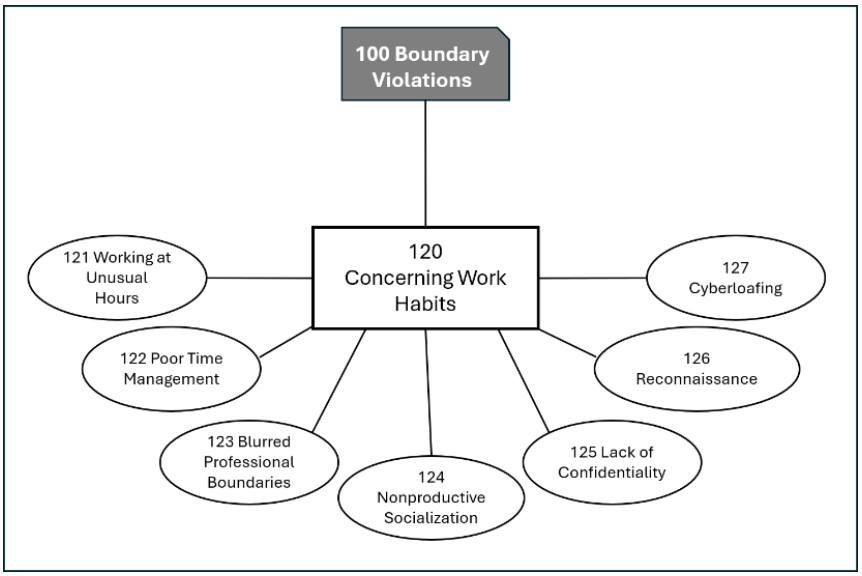

Figure 3. PRIs Associated with the SOFIT "Concerning Work Habits" Subclass

The specific makeup of the PRI list, the insider risk behaviors of interest, and the PRI probabilistic weights (which reflect the degree of association between PRIs and risk behaviors) are set in an expert knowledge elicitation/calibration process that is conducted with the organization's analysts. SOFIT, or any other desired PRI taxonomy, can provide the initial foundation for the deployed PRI knowledge base and threat behaviors, but the entire model is tailorable by the organization's insider risk management team to fit the organization's mission and priorities.

Figure 4 illustrates the hierarchical nature of risk propagation in a hypothetical sabotage scenario. In this example, interpersonal conflict, grievances, and other motivational cues combine with reconnaissance-like behavior and unauthorized operational actions. At lower levels of the model, these are represented as potential risk indicators; at higher levels, they are aggregated into patterns associated with candidate threat behaviors. The analyst can then inspect the strongest associations, review the supporting evidence, and determine whether the emerging pattern is most consistent with sabotage, negligence, or another explanation. In the figure, the blue circles represent the behavioral analytic risk propagation process that calculates the risk of the behavioral pattern comprising the observed combination of PRIs for an individual of interest. The nature of the calculation is also tailorable by the organization; a description of the mathematical model used in a Cogynt.ai insider risk management application, along with the results of a validation test of the model, is presented in [25].

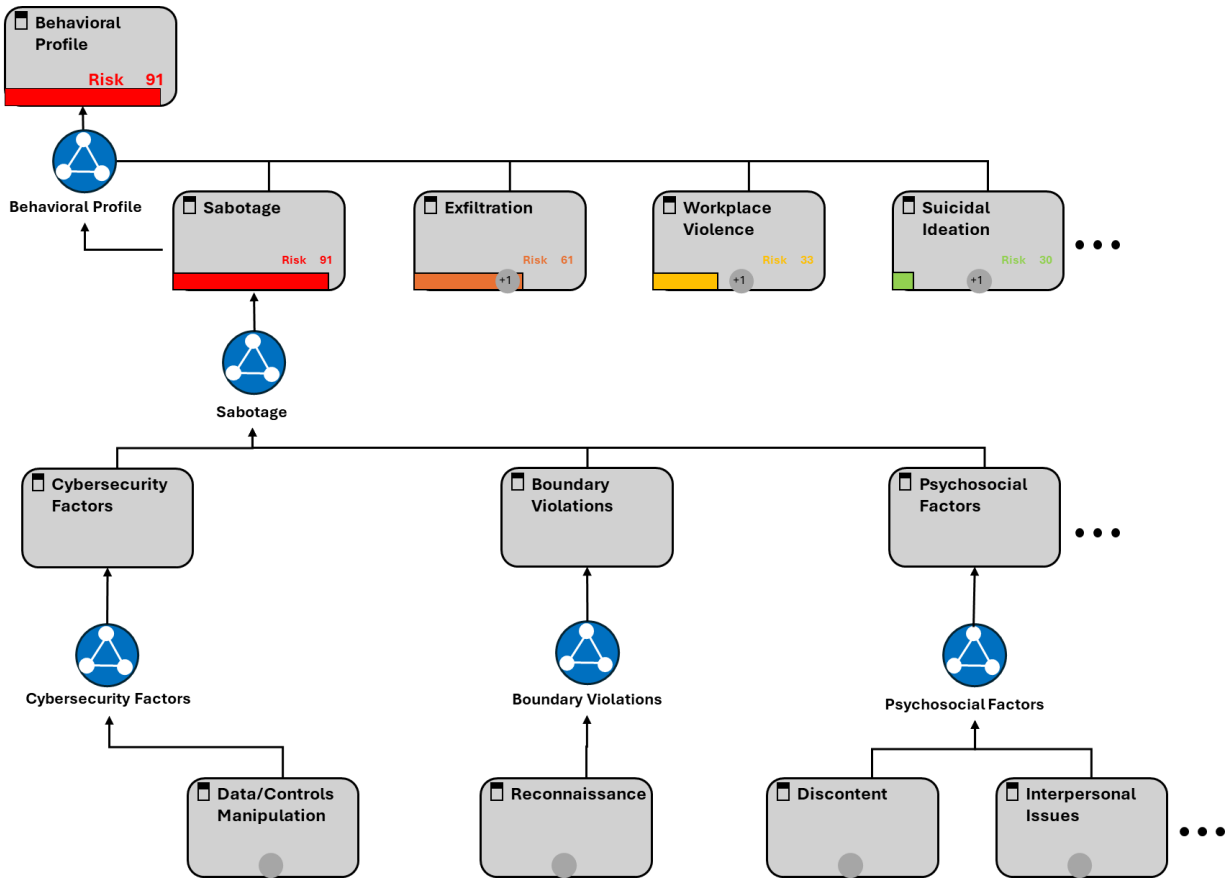

Figure 4. A Portion of the HCEP Diagram Showing the Mapping of Risk Indicators to Threat Behaviors in a Hypothetical Case of Sabotage

GOVERNANCE, IMPLEMENTATION, AND EVALUATION

Together, governance, implementation, and evaluation determine whether an insider risk capability can be deployed responsibly, trusted by stakeholders, and shown to provide operational value in the nuclear environment.

For deployment in nuclear environments, an insider risk analytic capability must be governed as carefully as the risks it is intended to detect. Programs should define clear authority for data access, model administration, case review, and intervention decisions; document lawful purposes and minimization rules for data use; and establish multidisciplinary oversight involving security, operations, HR, legal, privacy, and, where appropriate, safety stakeholders. These governance measures are necessary not only for compliance, but also for workforce legitimacy and program durability [1], [19]–[22].

Viewed through an organizational resilience lens, the value of this approach extends beyond improved detection. By combining transparent analytic reasoning with clear governance, multidisciplinary review, and proportionate intervention pathways, the capability can help organizations preserve trust, maintain continuity of operations, and strengthen decision quality under conditions of uncertainty. This framing may be especially important for audiences beyond analysts. Senior leaders, regulators, and operators are likely to view value not only in earlier warning, but also in the extent to which the program reinforces institutional trust, supports consistent and accountable decisions, and demonstrates that early indicators will be addressed in a lawful, fair, and operationally constructive manner.

A brief operational vignette helps illustrate the point. For example, if the model detects an emerging pattern involving unusual after-hours access to an MC&A database, repeated attempts to bypass two-person procedures for inventory adjustments, and recent signs of strain or conflict in a worker assigned to material balance area responsibilities, the initial response need not be punitive. Instead, the case can be routed for cross-functional review and a proportionate intervention such as supervisory engagement, temporary review of access privileges, retraining on MC&A procedures, or supportive personnel follow-up, depending on the evidence and context. In this way, the model supports earlier action that is both more targeted and less disruptive, reducing the likelihood that a manageable concern escalates into material control, safety, security, or workforce-trust problems.

In practice, trusted AI for insider risk should satisfy several conditions: the data sources and transformations should be documented; the inference logic should be reviewable; model performance should be periodically tested for validity, drift, and bias; and outputs should be framed as decision support rather than unchallengeable verdicts. Equally important, interventions should be proportionate to the level and quality of concern. Many cases will warrant supportive engagement, clarification, retraining, or temporary review rather than punitive escalation. This is particularly important to distinguish negligence, stress, and organizational friction from adversarial intent.

From an operational perspective, a practical implementation path for nuclear operators would begin with a constrained use case, such as privileged-user monitoring for MC&A systems, access anomalies in protected areas, or integrated review of sabotage-relevant precursor behaviors. Early pilot efforts should define data boundaries, establish review criteria, and measure analytic value against baseline practice. As described in the previous section, an early step in deploying an expert AI decision intelligence platform, such as Cogynt.ai, is to perform the detailed calibration and tailoring of the underlying PRI knowledge base and updating of the hierarchical pattern architecture and HCEP analysis to meet the requirements of the nuclear MC&A domain.

Once deployed, there are several validation options for testing or validating insider threat assessment tools [26] [27]. Useful evaluation metrics include reduction in analyst workload, alert precision, time to triage, proportion of alerts with sufficient evidentiary context, the ability to identify emerging cases earlier than conventional workflows, and several quantitative performance measures [27] such as false positive rates, accuracy, and signal detection metrics (e.g., area under the ROC curve, AUC). Ultimately, evaluation of the model's performance can be done using data gathered over time, preferably compared with historical baseline data. The challenge is to ensure that the same data outcomes are recorded under the two conditions. Evaluations of models based on real data are rare due to the lack of data and privacy constraints, but research studies using hypothetical or simulated/synthetic data have compared the outputs of risk models with expert judgments as proxies for ground truth (e.g., [24], [25], [28]).

One study [25] used hypothetical scenarios and expert judgments of insider risk to evaluate performance of several insider risk assessment modeling approaches, including the Cogynt.ai HCEP model. The pattern-based Cogynt.ai model outperformed others in predicting insider risks based on several measures including false positive rates, accuracy, and AUC. Another recent evaluation study applied directly to the nuclear sector compared two risk modeling approaches using data from more than 30,000 simulated operational technology/industrial control system scenarios [28]. A probabilistic modeling framework employing SOFIT behavioral PRIs and a contextual analysis that incremented the computed risk for certain combinations of early-stage indicators yielded earlier detection of insider activity compared to a baseline model, as determined by expert judgments. As suggested in earlier behavioral-analytic insider risk research (e.g., [24], [25]), the authors in [28] argue that conditioning cyber detection on identity-specific contextual risk signals can move detection earlier in the attack lifecycle, shifting risk assessment from a predominantly reactive posture to a more proactive one.

CONCLUSION

Insider risk in the nuclear sector is fundamentally a socio-technical problem that cannot be addressed through compliance checks, physical protection, or cybersecurity controls alone. Effective mitigation requires the ability to recognize early patterns of concern across human behavior, operational context, and cyber-physical activity, and to do so in a manner that is trustworthy, explainable, and operationally useful. The expert AI decision intelligence approach described in this paper offers one path toward that objective by combining behavioral science, heterogeneous data fusion, and auditable inference patterns to support whole-person risk assessment. If implemented with appropriate governance, privacy safeguards, and human oversight, such capabilities can help nuclear organizations move from reactive investigation toward earlier, more proportionate, and more resilient insider risk management—strengthening not only detection, but also institutional trust, continuity of operations, and the quality of decision-making under uncertainty.

REFERENCES

[1] International Atomic Energy Agency. (2020). *Preventive and Protective Measures Against Insider Threats* (Nuclear Security Series No. 8-G, Rev. 1). Vienna: IAEA.

[2] Bunn, M. (2023). Insider threats to nuclear security. In C. Hobbs, S. Tzinieris, & S. K. Aghara (Eds.), *The Oxford Handbook of Nuclear Security*. Oxford University Press.

[3] Cybersecurity Insiders. (2024). *2024 Insider Threat Report*. https://www.cybersecurity-insiders.com/portfolio/2024-insider-threat-report-gurucul/, accessed May 10, 2026.

[4] Nazir, S., Patel, S., & Patel, D. (2017). Assessing and augmenting SCADA cybersecurity: A survey of techniques. *Computers & Security, 70*, 436–454.

[5] Krause, T., Ernst, R., Klaer, B., Hacker, I., & Henze, M. (2021). Cybersecurity in power grids: Challenges and opportunities. *Sensors, 21*(18), 6225.

[6] Panful, B., Apaflo, B., Filani, A., Nnadi, K., & Hutchful, N. (2025). Human factor vulnerabilities in energy industry cybersecurity: Assessing employee awareness and behavior in breach prevention. *International Journal for Multidisciplinary Research, 7*(6).

[7] U.S. Atomic Energy Commission. (1961). *Report on the SL-1 Incident, January 3, 1961: The General Manager's Board of Investigation*. Idaho Falls, ID: Idaho Operations Office.

[8] General Electric Company Aircraft Nuclear Propulsion Department. (1962). *Final Report of SL-1 Recovery Operation, May 1961 Through July 1962* (IDO-19311). Idaho Falls, ID.

[9] Hammer, A. (2024). Mystery of America's first fatal nuclear disaster - with rumors still rife over 60 years later that explosion in remote Idaho town was triggered by one man's murderous rage amid LOVE TRIANGLE. Daily Mail, April 6, 2024. https://www.dailymail.co.uk/news/article-13271859/mystery-america-nuclear-disaster-idaho-love-triangle-murder-suicide.html, accessed May 10, 2026.

[10] U.S. Nuclear Regulatory Commission. (1979). *IE Circular No. 79-08: Attempted Extortion—Low Enriched Uranium*. https://www.nrc.gov/reading-rm/doc-collections/gen-comm/circulars/1979/cr79008.html, accessed May 10, 2026.

[11] Hendry, W.J. (1980). Safeguards System Response in an Actual Uranium Diversion Incident. *INMM Proceedings*. https://resources.inmm.org/annual-meeting-proceedings/safeguards-system-response-actual-uranium-diversion-incident, accessed May 10, 2026.

[12] PBS SoCal. (2012). *San Onofre Nuclear Plant: Possible "Tampering" of Generator Discovered*. https://www.pbssocal.org/redefine/san-onofre-nuclear-plant-possible-tampering-of-generator-discovered, accessed May 20, 2026.

[13] Zeller, T. (2012). San Onofre Nuclear Plant Investigating Possible Sabotage Of Safety System. *Huffington Post*, Nov 29, 2012. https://www.huffpost.com/entry/san-onofre-nuclear-plant-sabotage_n_2215260, accessed May 20, 2026.

[14] U.S. Department of Justice. (2016). *Former U.S. Nuclear Regulatory Commission Employee Sentenced to Prison for Attempted Spear-Phishing Attack on Department of Energy Computers*. https://www.justice.gov/archives/opa/pr/former-us-nuclear-regulatory-commission-employee-sentenced-prison-attempted-spear-phishing, accessed May 10, 2026.
[15] Center for Development of Security Excellence (CDSE). (n.d.). *Insider Threat Case Study: Charles H. Eccleston*. https://www.cdse.edu/Portals/124/Documents/casestudies/Insider-Threat-Case-Study-Charles-Eccleston.pdf, accessed May 10, 2026.
[16] World Nuclear News. (2014). *Doel 4 Shuts Down After Worker Action*. https://www.world-nuclear-news.org/Articles/Doel-4-shuts-down-after-worker-action, accessed May 10, 2026.
[17] VRT. (2019). Five years on, still no clue as to who sabotaged Doel nuclear reactor. https://www.vrt.be/vrtnws/en/2019/08/05/five-years-on-still-no-clue-as-to-who-sabotaged-doel-nuclear-rea/, accessed May 10, 2026.
[18] Brussels Times. (2021). Enquiry into nuclear plant sabotage comes to no conclusion. https://www.brusselstimes.com/181163/enquiry-into-nuclear-plant-sabotage-comes-to-no-conclusion, accessed May 10, 2026.
[19] Intelligence and National Security Alliance. (2025). *Recommendations when Using AI in Insider Risk Management*. Arlington, VA: INSA.
[20] World Institute for Nuclear Security. (2025). *Managing Insider Threats in the Nuclear Industry: A WINS International Best Practice Guide*. Vienna: WINS.
[21] National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1.
[22] International Atomic Energy Agency. (2021). *Enhancing Nuclear Security Culture in Organizations Associated with Nuclear and Other Radioactive Material* (Nuclear Security Series No. 38-T). Vienna: IAEA.
[23] Cogility Software. (2026). *Cogynt.ai Decision Intelligence Platform*.
[24] Greitzer, F. L., Purl, J., Leong, Y. M., & Becker, D. E. (2018). SOFIT: Sociotechnical and organizational factors for insider threat. In *IEEE Security and Privacy Workshops (SPW), Workshop on Research for Insider Threat (WRIT)* (pp. 197–206).
[25] Greitzer, F. L. (submitted for publication). From patterns to predictions: Insider risk modeling with a pattern-based behavioral analytic model. Prepublication copy available at: http://ssrn.com/abstract=5971319, accessed May 10, 2026.
[26] Intelligence and National Security Alliance. (2022). *Measuring the Effectiveness of Insider Threat Programs.* Arlington, VA: INSA, insa_wp_effectiveness.pdf, accessed May 20, 2026.
[27] Greitzer, F. L. & Ferryman, T. A. (2013). Methods and metrics for evaluating analytic insider threat tools. *IEEE Security and Privacy Workshop on Research for Insider Threat (WRIT). 10.1109/SPW.2013.34*, accessed May 20, 2026.
[28] Gabbay, L., Nickerson, C., Noonan, C. (2026). SHIFTING DETECTION LEFT: Identity-Conditioned Monitoring for Early Insider Threat Detection in High Consequence Systems. *International Conference on Nuclear Security (ICONS),* Vienna, Austria, May 11-15, 2026.

Frank L. Greitzer, PhD is Chief Behavioral Scientist at Cogility Software, where he supports development and deployment of advanced decision intelligence solutions, including insider risk management. With a PhD in mathematical psychology, he has worked as a research psychologist for the US Navy, a human factors and AI researcher in the aerospace industry, and Chief Scientist in cognitive informatics at the U.S. DOE Pacific Northwest National Laboratory (1992-2012). In 2012 he founded PsyberAnalytix to perform consulting in the intersection of social/behavioral sciences and information security. His contributions to research and practice include numerous journal articles, conference papers, and invited talks.