The Race Against Zero-Day Exploits: How Continuous Monitoring Can Shrink the Attack Window for the Kingdom's Critical Energy Infrastructure.

Abstract

A zero-day exploit gives the attacker a head start that patching alone cannot erase. In oil and gas facilities and power networks the gap is wider still, because operational technology is old, availability comes first and updates wait for planned shutdowns. This paper asks how continuous monitoring can shorten the attack window, defined here as the time between the first use of an unknown vulnerability and its effective containment, for critical energy infrastructure in the Kingdom of Saudi Arabia. A systematic integrative review of thirty refereed journal articles published between 2020 and 2025 was combined with a reading of two national control documents. Two studies anchor the synthesis: a comparative evaluation of artificial intelligence techniques for zero-day detection and a survey of anomaly detection in power grids. The synthesis yields a five-layer continuous exposure monitoring framework that links asset-aware network detection, security event correlation, threat intelligence, exploit-likelihood prioritisation and digital-twin testing to a set of response actions that do not depend on a vendor patch. An illustrative timeline built from invented values shows how earlier detection and compensating controls, rather than faster patching, account for most of the reduction in exposure. Evidence is supportive for anomaly-based detection and data-driven prioritisation, mixed for event correlation in operational networks and limited for digital twins and measured outcomes in Gulf energy settings. No primary or proprietary data were used.

1. Introduction

Energy infrastructure is a prize for capable attackers. Oil and gas production, refining, pipelines and electricity networks are run by industrial control systems that were designed for reliability and isolation rather than for hostile networks, and a successful intrusion can stop production, damage equipment or endanger staff. A survey of attacks on the oil and gas sector from the early 1990s to 2020 documented verified incidents in all three domains of the industry, upstream, midstream and downstream, and catalogued the attack paths used against them (Stergiopoulos et al., 2020). The Kingdom has direct experience of the threat: the Triton malware, which in 2017

UNDER PEER REVIEW IN IJAR

targeted the safety instrumentation of a Saudi petrochemical facility, could have caused catastrophic physical harm (Abraham et al., 2025).

The most dangerous intrusions use vulnerabilities that the defender does not yet know about. Such zero-day attacks tend to slip past signature-based and supervised detection because nothing in the training data resembles them (Zoppi et al., 2021). Once the vendor learns of the flaw, a patch must be written and released, and the release time itself varies with the characteristics of the vulnerability; zero-days that require privileges or affect confidentiality are less likely to be patched on time (Roumani, 2021). For an energy operator, the vendor's release date is only the start of a second delay, because updates to controllers and supervisory systems are tested and installed during scheduled outages. The period of exposure is therefore long, and most of it falls outside the operator's control.

That situation shifts attention from patching to monitoring. If an operator cannot close the vulnerability quickly, it can still try to see the attack early, limit what the attacker can reach and remove the foothold before damage is done. Continuous monitoring is the umbrella term for this capability. It covers network and host detection, the correlation of security events, the use of shared threat information and the ongoing assessment of which weaknesses are most likely to be exploited. Each strand has a substantial literature, but the strands are seldom brought together around the single question of how long an energy operator remains exposed.

This paper takes up that question. It is anchored in two studies. Ali et al. (2022) compared machine learning and deep learning techniques for detecting zero-day attacks and set out the challenges they face, and Madabhushi and Dewri (2023) structured a decade of research on anomaly and attack detection in power grids. Together they show what detection can achieve and where it falls short in the electricity sector. The present study extends their findings to the wider energy system and connects detection to prioritisation and response.

The aim is to explain how continuous monitoring can shorten the attack window for the Kingdom's critical energy infrastructure and to express that explanation as a practical framework. Three objectives follow: (O1) to identify the factors that lengthen exposure to zero-day exploitation in energy operational technology; (O2) to establish which monitoring capabilities can detect or contain such exploitation before a patch is deployed; and (O3) to assemble these capabilities into a framework and judge the strength of evidence behind each part. The

UNDER PEER REVIEW IN IJAP

corresponding research questions are: RQ1, what determines the length of the attack window in energy control environments? RQ2, which continuous monitoring capabilities can shorten it, and by what mechanism? RQ3, how strong is the evidence, and what conditions must hold for the framework to work in the Saudi context?

After this introduction, the paper reviews background on exposure and monitoring, explains the review method, reports findings against the three questions, presents the framework with an illustrative timeline, discusses evidence strength and implications, and closes with conclusions.

2. Background

2.1 The anatomy of an attack window

The life of a vulnerability can be traced through a sequence of events: the flaw is introduced, discovered by someone, exploited, disclosed, patched by the vendor and finally remediated by the user. A zero-day is a flaw that is exploited before the vendor or defender knows of it. The defender's exposure begins with the first exploitation and ends only when the attacker can no longer use the flaw to cause harm, whether because the patch is installed, the vulnerable service is isolated or the intruder is removed. In this paper that interval is called the attack window.

Two features of the window matter for design. The first is that the vendor's share of the delay is outside the operator's control and depends on the vulnerability itself (Roumani, 2021). The second is that, after disclosure, operators face far more vulnerabilities than they can fix at once. Firms commonly prioritise with crude heuristics and are still breached through known flaws whose fixes have been available for a long time (Jacobs et al., 2020). Shortening the window therefore depends as much on early sight of exploitation and on choosing the right fixes as on the speed of any single patch.

2.2 Why energy control systems stay exposed for longer

Energy control environments combine age, connectivity and intolerance of downtime. Critical infrastructure operators are moving modern information technology into monolithic operational technology architectures, yet the complexity and age of these systems prevent defenders from applying contemporary controls in a timely manner (Makrakis et al., 2021). Energy-sector supervisory and control systems follow different standards from ordinary information systems, and their availability is of paramount importance (Alghassab, 2021). Legacy equipment will

UNDER PEER REVIEW IN IJAP

continue to carry vital infrastructure for years, which is why threat modelling of legacy industrial systems in an oil and gas setting combined system-, attacker- and risk-centred methods to expose the full attack surface (Badawy et al., 2024).

Offshore and remote assets add further exposure. Offshore oil production depends on remote access to and control of complex hardware, achieved by integrating supervisory control and data acquisition with industrial Internet of Things technologies, and an analysis of a subsea control architecture identified vulnerabilities across its components (Mohammed et al., 2022). A systematic review of digital transformation in the energy sector found cybersecurity challenges across oil and gas, electricity and nuclear organisations, and called for more research in this area (Saeed et al., 2024). In power networks, the consequences can propagate: a study of historic blackouts and simulation showed how cyber attacks can activate cyber-physical factors that cause and accelerate cascading failures (Rajkumar et al., 2023).

The adversaries who matter most are patient. Advanced persistent threats are stealthy, multistage campaigns tailored to a target and often remain undetected until their consequences are evident (Abraham et al., 2025). Ransomware is a further concern, and its prediction in industrial systems is hampered by proprietary software and limited visibility into operations (Gazzan & Sheldon, 2023). Both observations point to the same weakness: the defender often cannot see what is happening inside the control network.

2.3 Monitoring as a counterweight

Monitoring research for industrial and grid settings has grown quickly. Network-based intrusion detection for supervisory control systems must respect real-time constraints, distinctive traffic features and application-layer protocols, and a review of such systems recorded considerable progress in machine-learning detection methods and testbeds (Rakas et al., 2020). Surveys of machine learning detection in smart grids examined where such detectors are applied across transmission and distribution, how datasets are generated and how detectors are evaluated (Sahani et al., 2023). At the level of the security operation, information and event management platforms collect, normalise and correlate events to give wide visibility and to reduce response time (González-Granadillo et al., 2021). Each of these is a partial answer to the visibility problem described above.

UNDER PEER REVIEW IN IJAP

Grid research also offers defences that act on the physical layer. A review of cyber-physical attack and defence in the smart grid proposed a unified state-space model in which attacks and defences can be generalised, and discussed operational and informational defences such as moving target defence, watermarking and data-driven methods (Zhang et al., 2021). Such techniques matter for zero-days because they can make an attack harder to execute or easier to notice even when the exploited flaw is unknown. They complement, rather than replace, network monitoring, since they depend on detailed models of the physical process that only the operator holds.

2.4 The Saudi regulatory context

Two national documents set expectations for operators in the Kingdom (Table 1). The Essential Cybersecurity Controls of the National Cybersecurity Authority apply to government entities and to private entities that own, operate or host critical national infrastructure, and include subdomains for vulnerability management, event logs and monitoring management, and incident and threat management (Table 1, D1). The vulnerability management subdomain requires, at a minimum, periodic assessment and detection, classification by severity, remediation according to classification and risk, verified patching and subscription to trusted sources of vulnerability information. The Operational Technology Cybersecurity Controls, published in 2022 as an extension of the essential controls, apply to industrial control systems in critical facilities and contain four domains, 23 subdomains, 47 main controls and 122 subcontrols, with the controls required depending on the criticality level of the facility (Table 1, D2). These documents define obligations; they do not specify how monitoring should be arranged to shorten exposure, which is the gap this paper addresses.

The two documents also differ in emphasis. The essential controls treat vulnerability management largely as a cycle of assessment and remediation, while the operational technology controls recognise that industrial systems must be protected according to the consequence of their disruption. A monitoring framework for energy operators has to satisfy both: it must support the required cycle while accepting that, in control networks, remediation is often deferred and protection must come from elsewhere in the meantime.

UNDER PEER REVIEW IN IJAR

Table 1. National documentary sources consulted (not counted among the thirty reviewed articles)

Code	Document	Issuer and year	Provisions relevant to this study
D1	Essential Cybersecurity Controls (ECC-2:2024)	National Cybersecurity Authority, 2024	Applies to government entities and private owners or operators of critical national infrastructure; subdomains on vulnerability management (2-10), event logs and monitoring management (2-12) and incident and threat management (2-13); minimum vulnerability requirements include periodic assessment, severity classification, risk-based remediation, verified patching and trusted vulnerability sources
D2	Operational Technology Cybersecurity Controls (OTCC-1:2022)	National Cybersecurity Authority, 2022	Extension of the essential controls for industrial control systems in critical facilities; four domains, 23 subdomains, 47 main controls and 122 subcontrols; required controls depend on the criticality level of the facility

3. Method

3.1 Review design

The evidence is heterogeneous, which is why a systematic integrative review was used. It includes detection experiments, surveys, design studies, case analyses and conceptual frameworks, and no single outcome measure is shared across them, so statistical pooling would be inappropriate. An integrative review allows such evidence to be combined into a coherent explanation, while systematic procedures for searching, screening and coding make the process traceable. The output is an explanatory framework rather than an effect estimate.

3.2 Search and selection

In November 2026 the author queried an open catalogue of scholarly literature, and the registration records held by publishers were used to confirm identifiers, publication years and the absence of retraction notices. Queries combined terms for zero-day attacks, exploit prediction, patch release and prioritisation, intrusion and anomaly detection, security operations and event management, threat intelligence, digital twins and industrial control systems with terms for oil and gas, power grids and energy infrastructure. Inclusion required a refereed journal article published between 2020 and 2025, a focus relevant to detecting, prioritising or containing exploitation in critical or industrial infrastructure, and an abstract or full text sufficient for coding. Conference papers, preprints and items with unresolved identifiers were set aside, as

UNDER PEER REVIEW IN IJAR

were articles concerned only with consumer devices. Screening produced thirty articles. The two anchor studies (Ali et al., 2022; Madabhushi & Dewri, 2023) were chosen first because they summarise detection capability for zero-days and for power grids respectively.

3.3 Coding and synthesis

For every article the coding recorded the stage of the attack window it addresses (before disclosure, between disclosure and patch, or after patch release), the capability it describes, the setting, the type of evidence and the main finding. Capabilities were then grouped into layers, and each layer was judged supportive, mixed or limited on the basis of the volume and consistency of evidence and how closely the settings match energy operations. National control documents were read for their requirements and are listed separately in Table 1; they are not counted among the thirty articles. The illustrative timeline in Section 5.3 uses invented values chosen by the author to show the logic of the framework. It does not describe any real organisation, incident or facility, and the study involved no primary or proprietary data.

4. Findings

4.1 What determines the length of the attack window (RQ1)

The review points to four determinants. Vendor delay, already described, comes first; it is shaped by attack vector, required privileges, scope and impact (Roumani, 2021). The second is the operator's remediation delay. In energy control environments, controllers and supervisory servers cannot be restarted at will, so remediation waits for maintenance windows, and legacy components may never receive a fix (Badawy et al., 2024; Makrakis et al., 2021).

The third determinant is detection delay, the time before the defender notices exploitation. This is where the window is longest and least visible. Persistent attackers deliberately move slowly (Abraham et al., 2025), signature and supervised methods tend to miss novel attacks (Zoppi et al., 2021), and the limited visibility of proprietary industrial systems hinders early warning (Gazzan & Sheldon, 2023). The fourth determinant is prioritisation error. When remediation is guided by severity scores and expert opinion alone, effort is spread across many low-risk flaws while high-risk ones remain open (Jacobs et al., 2020, 2021).

Of the four, only the first lies wholly outside the operator's influence. The other three can be shortened by monitoring, and the detection delay offers the largest gain because it can run for

UNDER PEER REVIEW IN IJAR

months while the other stages are measured in days or weeks. In short, the attack window in an energy setting is long mainly because exploitation goes unseen, not because patches are slow.

4.2 How continuous monitoring shortens the window (RQ2)

Five capabilities emerge from the evidence. The first is anomaly-based detection in the operational network. Ali et al. (2022) reviewed machine and deep learning algorithms for zero-day detection and compared reported accuracy, precision and recall, and found that the techniques still face challenges in detecting such attacks. Unsupervised algorithms show potential because they model normal behaviour rather than known attacks, although quantitative support for evaluating them remains scarce and the choice of features and metrics strongly affects results (Zoppi et al., 2021). In power grids, anomaly detection research of the past decade spans many methods and data sources, with the grid's many entry points making detection central to defence (Madabhushi & Dewri, 2023). A concatenated convolutional and recurrent network deployed against a benchmark dataset for smart metering reported 98.84% detection accuracy (Diaba et al., 2022), but benchmark results of this kind should be read as upper bounds rather than field performance. The value of testbeds and datasets for training and testing detectors in industrial settings has been documented in a dedicated survey (Conti et al., 2021).

The second capability is correlation of events across sources. Detection signals from network sensors, hosts and controllers are of limited use unless they reach analysts in context. A survey of event management platforms in critical infrastructure noted their growing convergence with big data analytics and listed enhancements needed for the next generation (González-Granadillo et al., 2021). A purpose-built platform for the smart grid showed that conventional tools ignore grid-specific features such as legacy protocols, and demonstrated detection, normalisation and correlation of attacks tailored to those protocols (Radoglou-Grammatikis et al., 2021). Correlation turns individual anomalies into incidents and is therefore the step at which detection delay ends.

The third capability is threat intelligence. Indicators of compromise allow responders to detect the triggers of malicious activity and to act against a comparable intrusion sooner, but their use in industrial environments requires an understanding of context and representation standards (Asiri et al., 2023). A framework for continuous control of access rights in an industrial setting used incident information shared by trusted entities to update security policies and prevent

UNDER PEER REVIEW IN IJAR

incidents (Osliak et al., 2023). Intelligence matters for zero-days because exploitation of a new flaw at one operator often becomes known before a patch exists, and shared indicators can then shorten detection at others.

The fourth capability is exploit-likelihood prioritisation once a flaw is disclosed. The Exploit Prediction Scoring System estimates how likely a flaw is to see exploitation in the wild during the year after it becomes public and reported a receiver operating characteristic area under the curve of 0.838 (Jacobs et al., 2021). Remediation strategies built with machine learning on multiple datasets, using exploitation observed in real firm networks as the outcome, were compared on how well they balance coverage against efficiency (Jacobs et al., 2020). Context matters as well. An automated, context-aware approach to vulnerability risk management let the organisation choose a mode and weight the criteria used to rank patches (Ahmadi Mehri et al., 2022). Prediction is also possible early. Models trained on the initial vulnerability record and linked discussions can predict exploitability soon after disclosure, although pre-trained language models did not perform as expected without further security training (Iannone et al., 2024).

The fifth capability is safe testing through replicas. Operators resist actions that imply downtime, and digital twins built from real-time, high-fidelity replicas of programmable logic controllers can provide an environment for security analysis and for validating machine learning models against threats such as denial of service (Sousa et al., 2021). For zero-day response, a twin allows compensating controls and emergency fixes to be tested without touching production.

These capabilities shorten different parts of the window. Detection and correlation cut the time to discovery; intelligence cuts it further and supports containment; prioritisation reduces the time that the most dangerous disclosed flaws stay open; and replicas reduce the time needed to validate a change. None requires the vendor patch to arrive sooner.

The capabilities also depend on one another. Anomaly detectors produce many false alarms when baselines are poor, and correlation with asset context is what allows an analyst to discard benign deviations quickly. Intelligence is only useful if local telemetry can be searched for the shared indicators, which presupposes adequate logging. Prioritisation needs an accurate inventory of where each vulnerable product is installed, and replica testing needs current configurations of the controllers being copied. A monitoring programme that buys sensors without building these links is likely to add noise rather than shorten exposure. This

UNDER PEER REVIEW IN IJAR

interdependence is the reason for presenting the capabilities as layers of one framework rather than as a list of separate tools.

4.3 Evidence strength and enabling conditions (RQ3)

The strength of evidence varies across the five capabilities (Table 2). Support for anomaly-based detection is the broadest, with several reviews and experiments in industrial and grid settings, but almost all of it comes from benchmark datasets and laboratory testbeds (Conti et al., 2021; Rakas et al., 2020; Sahani et al., 2023). Exploit-likelihood prioritisation is supported by large-scale empirical work, though in information technology rather than control systems (Jacobs et al., 2020, 2021). Evidence on event correlation in operational networks is mixed: the capability is well described, yet grid-specific platforms remain few (Radoglou-Grammatikis et al., 2021). Threat intelligence for industrial settings is mixed, because standards and tools exist but their transfer to industrial contexts is still being worked out (Asiri et al., 2023). Digital twins for security are at the demonstration stage (Sousa et al., 2021).

Three enabling conditions recur. People come first: security operations centres depend on the link between human analysts and technology through defined processes, an area that academic work has neglected (Vielberth et al., 2020), and a systematic method for measuring analyst performance was validated in two centres (Agyepong et al., 2023). Training is the second. A review of training for critical infrastructure protection, covering energy among other sectors, found that hands-on, team-based and realistic training is preferred but that agreement on optimal methods is lacking (Chowdhury & Gkioulos, 2021). The third condition is knowledge of the system itself, since detection and threat modelling both assume that assets and communication paths are known (Badawy et al., 2024). Where these conditions are absent, monitoring tools generate alerts without shortening exposure.

Table 2. Monitoring layers, their effect on the attack window and the strength of supporting evidence

Layer	Main effect on the attack window	Key sources	Evidence
1. Asset-aware network detection	Shortens detection delay by flagging deviations from normal behaviour, including unknown attacks	Ali et al. (2022); Madabhushi & Dewri (2023); Zoppi et al. (2021); Diaba et al. (2022); Sahani et al. (2023)	Supportive (mostly benchmark data)
2. Event correlation	Ends detection delay by turning alarms into incidents with asset	González-Granadillo et al. (2021); Radoglou-Grammatikis et al. (2021);	Mixed

UNDER PEER REVIEW IN IJAR

Layer	Main effect on the attack window	Key sources	Evidence
	context	Vielberth et al. (2020)	
3. Threat intelligence	Brings forward detection and containment through shared indicators	Asiri et al. (2023); Osliak et al. (2023)	Mixed
4. Exploit-likelihood prioritisation	Reduces time that the most dangerous disclosed flaws remain open	Jacobs et al. (2020, 2021); Ahmadi Mehri et al. (2022); Iannone et al. (2024)	Supportive (enterprise IT)
5. Replica testing	Shortens validation of containment steps and emergency changes	Sousa et al. (2021); Conti et al. (2021)	Limited
Enabling conditions	Determine whether signals become actions	Agyepong et al. (2023); Chowdhury & Gkioulos (2021); Badawy et al. (2024)	Mixed

5. A Continuous Exposure Monitoring Framework

5.1 Structure

The framework arranges the five capabilities as layers that feed a common decision point (Figure 1). Layer one, asset-aware network detection, establishes behavioural baselines for each control network segment and flags deviations, using unsupervised methods where attacks are unknown. Layer two, event correlation, joins network, host and controller events with asset context so that analysts see incidents rather than isolated alarms. The third layer, threat intelligence, imports indicators and attack descriptions from trusted partners and national bodies and exports local findings. Exploit-likelihood prioritisation forms layer four and ranks disclosed vulnerabilities by predicted exploitation and by the criticality of the affected facility. Finally, replica testing uses a digital twin or a representative test bench to validate containment steps and emergency changes.

The decision point chooses among response actions that do not wait for a vendor patch: segmenting or isolating an affected zone, disabling an exposed service or remote access path, adding temporary detection rules, increasing logging on the affected assets and, when available, applying a verified patch in the next safe window. The framework is continuous in the sense that each layer runs all the time and each response feeds back into baselines, rules and priorities.

UNDER PEER REVIEW IN IJAR

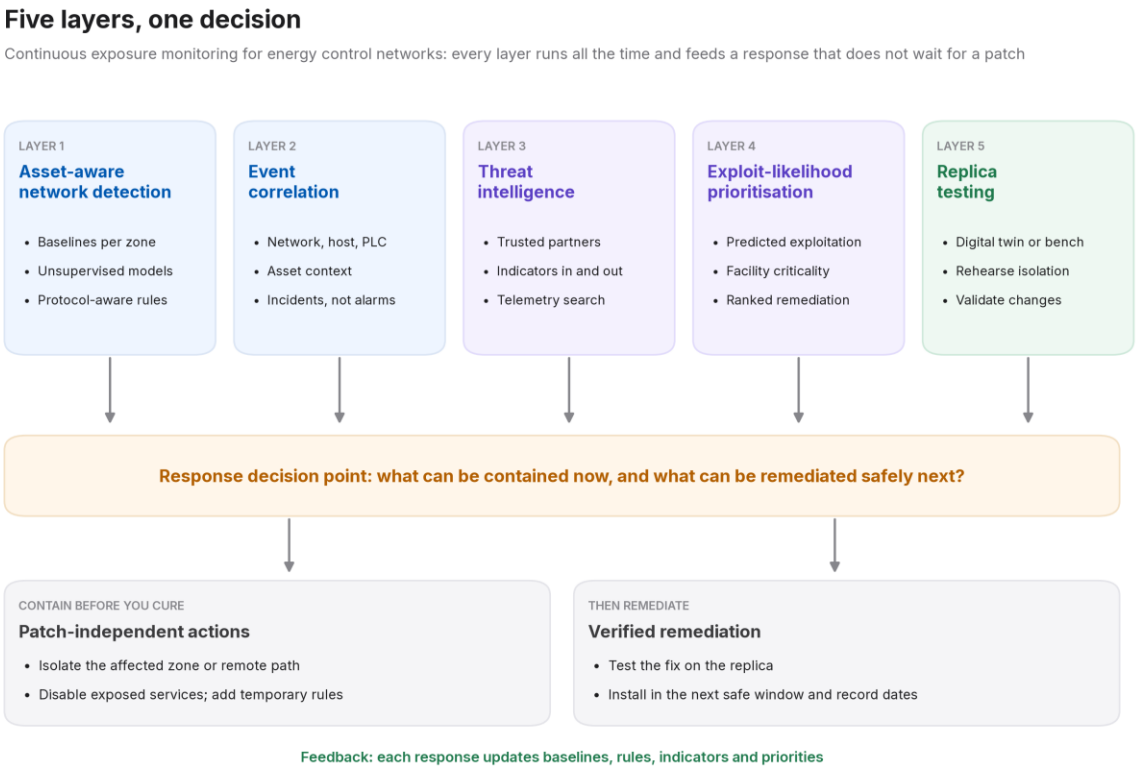


Figure 1. Continuous exposure monitoring framework: five monitoring layers feeding a patch-independent response decision

5.2 Operating principles

Four principles guide use. First, measure the window. Operators should record for each significant vulnerability or incident the dates of first evidence of exploitation, detection, containment and remediation, because the framework's purpose is to shorten the interval between the first and the third. Second, contain before you cure. The evidence on release and remediation delays shows that waiting for a patch is not a strategy in control environments, so containment should be planned in advance for each critical zone (Roumani, 2021; Makrakis et al., 2021).

Third, rank by likelihood and consequence together. Severity scores alone are a weak guide; predicted exploitation and the consequence of failure at the affected facility should both inform the order of remediation (Ahmadi Mehri et al., 2022; Jacobs et al., 2020). Rajkumar et al. (2023) showed how cyber attacks can trigger cascading failures, so facility criticality should be judged at system level. Fourth, treat the operating centre as part of the control system. Analyst

UNDER PEER RE

processes, performance measures and training determine whether signals become actions (Agyepong et al., 2023; Vielberth et al., 2020).

These principles align with the national controls. Periodic assessment, classification, risk-based remediation and subscription to trusted sources are already required, and event logging and monitoring form their own subdomain (Table 1, D1). The framework adds an explicit target, the length of the attack window, and the means of shortening it.

5.3 Illustrative attack-window timeline

Figure 2 shows how the framework changes exposure, using an invented scenario. All values are illustrative and chosen by the author; they do not describe any real incident, facility or operator. A hypothetical zero-day affecting an engineering workstation in a gas processing control network is first exploited on day 0. Under periodic assessment, with quarterly scanning and monthly log review, the intrusion is noticed on day 41 and contained on day 48, while the vendor patch, released on day 30, is installed during a planned outage on day 95. Under continuous monitoring, an anomaly in workstation traffic is flagged on day 3, correlated with a partner's indicator and contained on day 5 by isolating the workstation segment and blocking the remote path, and the patch is validated on a replica and installed on day 60.

In this example the attack window, measured from first exploitation to containment, falls from 48 days to 5. The time to remediation also falls, from 95 to 60 days, but that change is smaller and matters less, because the attacker lost access once the segment was isolated. The example therefore illustrates the central argument: in energy control environments, most of the reduction in exposure comes from earlier detection and planned containment rather than faster patching. Real values will vary widely, and the example should not be read as an estimate of any operator's performance.

UNDER PEER REVIEW IN IJAR

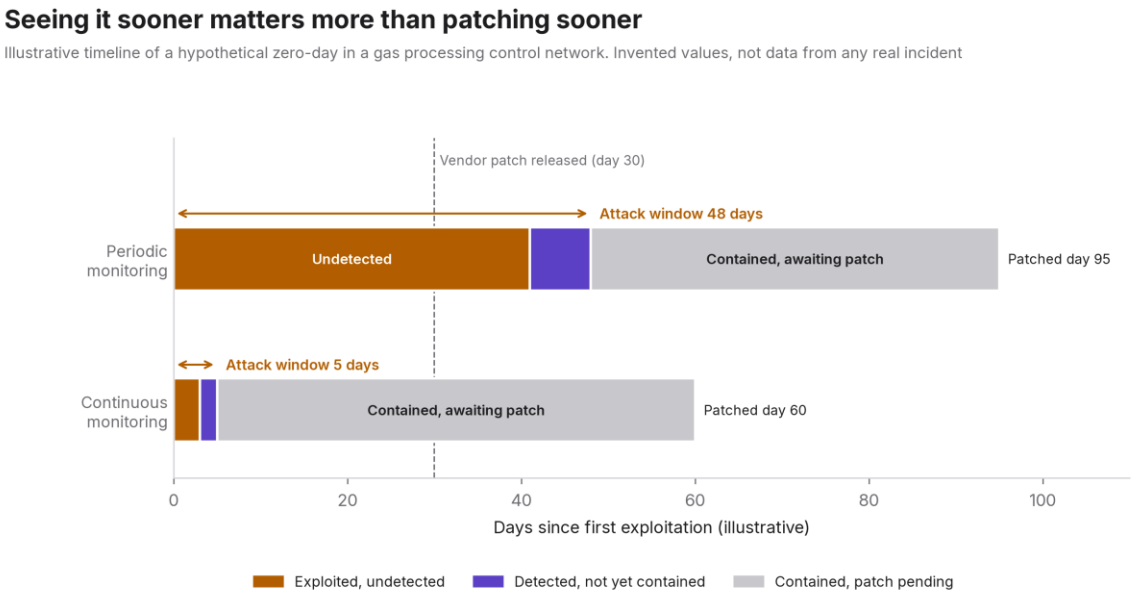


Figure 2. Illustrative attack-window timeline under periodic and continuous monitoring (invented values, not data from any real incident)

6. Discussion

6.1 What the review adds

The two anchor studies speak mainly to detection. Ali et al. (2022) showed that artificial intelligence techniques for zero-day detection still face open challenges, and Madabhushi and Dewri (2023) showed how varied anomaly detection research for power grids has become. This review places detection within a timeline and shows that it is one of several levers. Its contribution is to connect detection research with work on exploit prediction, operations management, intelligence sharing and safe testing, and to argue that the right outcome measure is the length of exposure rather than detection accuracy alone.

The evidence also exposes a mismatch. The most rigorous quantitative work on prioritisation comes from enterprise information technology (Jacobs et al., 2021; Iannone et al., 2024), while energy-specific studies are mostly surveys, design proposals and testbed experiments (Saeed et al., 2024; Mohammed et al., 2022). Few studies report how long real energy operators remain exposed or how monitoring changes that time. The framework is therefore a structured hypothesis that energy operators and researchers can test, not a proven solution.

UNDER PEER REVIEW

6.2 Implications for operators in the Kingdom

For Saudi energy operators, three implications stand out. One is to make detection in control networks a priority investment rather than an extension of office network monitoring. Grid- and plant-specific protocols, timing constraints and legacy devices require tailored sensors and rules (Rakas et al., 2020; Radoglou-Grammatikis et al., 2021). Another is to plan containment zone by zone before an incident occurs, using replicas or test benches to rehearse isolation steps, so that the decision point can act within hours. A third is for operators and national bodies to strengthen the sharing of indicators among energy entities, because shared indicators convert one victim's discovery into early warning for others (Asiri et al., 2023; Osliak et al., 2023).

Workforce capacity is the binding constraint behind all three. Monitoring at this depth needs analysts who understand both security and process control, and the training literature suggests that realistic, hands-on exercises are the preferred route (Chowdhury & Gkioulos, 2021). The facility levels used in the operational technology controls offer a natural basis for deciding where the full framework should apply first (Table 1, D2).

6.3 Limitations and further research

The review has limits. It is restricted to thirty journal articles in English from a single six-year period, and relevant evidence in conference papers, vendor reports and non-English sources was not included. Most cited studies use laboratory datasets, so detection performance in the field is likely to be lower than reported. The national documents were read only for their stated requirements. The illustrative timeline is a teaching device, not evidence.

Further research should measure attack windows in real energy settings, compare periodic and continuous arrangements on the same assets, adapt exploit prediction to industrial vulnerabilities and test whether digital twins shorten the time needed to validate containment. Studies of analyst workload and training in Gulf operating centres would also help to show whether the human conditions identified here can be met at scale.

7. Conclusion

Zero-day exploits will continue to reach energy control systems before patches do. The review shows that the attack window in such environments is long mainly because exploitation goes unseen and remediation must wait for safe windows, and that continuous monitoring can shorten

UNDER PEER REVIEW IN IJAR

it by detecting anomalies, correlating events, importing intelligence, prioritising by likelihood and testing containment on replicas. In answer to the research questions, the window is driven by vendor, detection, prioritisation and remediation delays (RQ1); the five monitoring layers act mainly on detection and containment (RQ2); and the evidence is supportive for anomaly detection and prioritisation but mixed or limited elsewhere, with skilled people, training and asset knowledge as preconditions (RQ3). For the Kingdom's critical energy infrastructure, the practical message is to measure exposure, plan containment in advance and treat monitoring as a core control rather than an add-on.

UNDER PEER REVIEW IN IJAR

References

1. Abraham, D., Houmb, S. H., & Erdodi, L. (2025). Cyber-attacks on energy infrastructure—a literature overview and perspectives on the current situation. Applied Sciences, 15(17), Article 9233. https://doi.org/10.3390/app15179233

2. Agyepong, E., Cherdantseva, Y., Reinecke, P., & Burnap, P. (2023). A systematic method for measuring the performance of a cyber security operations centre analyst. Computers & Security, 124, Article 102959. https://doi.org/10.1016/j.cose.2022.102959

3. Ahmadi Mehri, V., Arlos, P., & Casalicchio, E. (2022). Automated context-aware vulnerability risk management for patch prioritization. Electronics, 11(21), Article 3580. https://doi.org/10.3390/electronics11213580

4. Alghassab, M. (2021). Analyzing the impact of cybersecurity on monitoring and control systems in the energy sector. Energies, 15(1), Article 218. https://doi.org/10.3390/en15010218

5. Ali, S., Rehman, S. U., Imran, A., Adeem, G., Iqbal, Z., & Kim, K.-I. (2022). Comparative evaluation of AI-based techniques for zero-day attacks detection. Electronics, 11(23), Article 3934. https://doi.org/10.3390/electronics11233934

6. Asiri, M., Saxena, N., Gjomemo, R., & Burnap, P. (2023). Understanding indicators of compromise against cyber-attacks in industrial control systems: A security perspective. ACM Transactions on Cyber-Physical Systems, 7(2), 1–33. https://doi.org/10.1145/3587255

7. Badawy, M., Sherief, N. H., & Abdel-Hamid, A. A. (2024). Legacy ICS cybersecurity assessment using hybrid threat modeling—an oil and gas sector case study. Applied Sciences, 14(18), Article 8398. https://doi.org/10.3390/app14188398

8. Chowdhury, N., & Gkioulos, V. (2021). Cyber security training for critical infrastructure protection: A literature review. Computer Science Review, 40, Article 100361. https://doi.org/10.1016/j.cosrev.2021.100361

9. Conti, M., Donadel, D., & Turrin, F. (2021). A survey on industrial control system testbeds and datasets for security research. IEEE Communications Surveys & Tutorials, 23(4), 2248–2294. https://doi.org/10.1109/comst.2021.3094360

UNDER PEER REVIEW IN IJAR

10. Diaba, S. Y., Shafie-khah, M., & Elmusrati, M. (2022). On the performance metrics for cyber-physical attack detection in smart grid. Soft Computing, 26(23), 13109–13118. https://doi.org/10.1007/s00500-022-06761-1

11. Gazzan, M., & Sheldon, F. T. (2023). Opportunities for early detection and prediction of ransomware attacks against industrial control systems. Future Internet, 15(4), Article 144. https://doi.org/10.3390/fi15040144

12. González-Granadillo, G., González-Zarzosa, S., & Diaz, R. (2021). Security information and event management (SIEM): Analysis, trends, and usage in critical infrastructures. Sensors, 21(14), Article 4759. https://doi.org/10.3390/s21144759

13. Iannone, E., Sellitto, G., Iaccarino, E., Ferrucci, F., De Lucia, A., & Palomba, F. (2024). Early and realistic exploitability prediction of just-disclosed software vulnerabilities: How reliable can it be? ACM Transactions on Software Engineering and Methodology, 33(6), 1–41. https://doi.org/10.1145/3654443

14. Jacobs, J., Romanosky, S., Adjerid, I., & Baker, W. (2020). Improving vulnerability remediation through better exploit prediction. Journal of Cybersecurity, 6(1), Article tyaa015. https://doi.org/10.1093/cybsec/tyaa015

15. Jacobs, J., Romanosky, S., Edwards, B., Adjerid, I., & Roytman, M. (2021). Exploit Prediction Scoring System (EPSS). Digital Threats: Research and Practice, 2(3), 1–17. https://doi.org/10.1145/3436242

16. Madabhushi, S., & Dewri, R. (2023). A survey of anomaly detection methods for power grids. International Journal of Information Security, 22(6), 1799–1832. https://doi.org/10.1007/s10207-023-00720-z

17. Makrakis, G. M., Kolias, C., Kambourakis, G., Rieger, C., & Benjamin, J. (2021). Industrial and critical infrastructure security: Technical analysis of real-life security incidents. IEEE Access, 9, 165295–165325. https://doi.org/10.1109/access.2021.3133348

18. Mohammed, A. S., Reinecke, P., Burnap, P., Rana, O., & Anthi, E. (2022). Cybersecurity challenges in the offshore oil and gas industry: An industrial cyber-physical systems (ICPS) perspective. ACM Transactions on Cyber-Physical Systems, 6(3), 1–27. https://doi.org/10.1145/3548691

UNDER PEER REVIEW IN IJAR

19. Osliak, O., Saracino, A., Martinelli, F., & Mori, P. (2023). Cyber threat intelligence for critical infrastructure security. Concurrency and Computation: Practice and Experience, 35(23), Article e7759. https://doi.org/10.1002/cpe.7759

20. Radoglou-Grammatikis, P., Sarigiannidis, P., Iturbe, E., Rios, E., Martinez, S., Sarigiannidis, A., Eftathopoulos, G., Spyridis, Y., Sesis, A., Vakakis, N., Tzovaras, D., Kafetzakis, E., Giannoulakis, I., Tzifas, M., Giannakoulias, A., Angelopoulos, M., & Ramos, F. (2021). SPEAR SIEM: A security information and event management system for the smart grid. Computer Networks, 193, Article 108008. https://doi.org/10.1016/j.comnet.2021.108008

21. Rajkumar, V. S., Ştefanov, A., Presekal, A., Palensky, P., & Torres, J. L. R. (2023). Cyber attacks on power grids: Causes and propagation of cascading failures. IEEE Access, 11, 103154–103176. https://doi.org/10.1109/access.2023.3317695

22. Rakas, S. V. B., Stojanovic, M. D., & Markovic-Petrovic, J. D. (2020). A review of research work on network-based SCADA intrusion detection systems. IEEE Access, 8, 93083–93108. https://doi.org/10.1109/access.2020.2994961

23. Roumani, Y. (2021). Patching zero-day vulnerabilities: An empirical analysis. Journal of Cybersecurity, 7(1), Article tyab023. https://doi.org/10.1093/cybsec/tyab023

24. Saeed, S., Gull, H., Aldossary, M. M., Altamimi, A. F., Alshahrani, M. S., Saqib, M., Zafar Iqbal, S., & Almuhaideb, A. M. (2024). Digital transformation in energy sector: Cybersecurity challenges and implications. Information, 15(12), Article 764. https://doi.org/10.3390/info15120764

25. Sahani, N., Zhu, R., Cho, J.-H., & Liu, C.-C. (2023). Machine learning-based intrusion detection for smart grid computing: A survey. ACM Transactions on Cyber-Physical Systems, 7(2), 1–31. https://doi.org/10.1145/3578366

26. Sousa, B., Arieiro, M., Pereira, V., Correia, J., Lourenço, N., & Cruz, T. (2021). ELEGANT: Security of critical infrastructures with digital twins. IEEE Access, 9, 107574–107588. https://doi.org/10.1109/access.2021.3100708

UNDER PEER REVIEW IN IJAR

27. Stergiopoulos, G., Gritzalis, D. A., & Limnaios, E. (2020). Cyber-attacks on the oil & gas sector: A survey on incident assessment and attack patterns. IEEE Access, 8, 128440–128475. https://doi.org/10.1109/access.2020.3007960

28. Vielberth, M., Bohm, F., Fichtinger, I., & Pernul, G. (2020). Security operations center: A systematic study and open challenges. IEEE Access, 8, 227756–227779. https://doi.org/10.1109/access.2020.3045514

29. Zhang, H., Liu, B., & Wu, H. (2021). Smart grid cyber-physical attack and defense: A review. IEEE Access, 9, 29641–29659. https://doi.org/10.1109/access.2021.3058628

30. Zoppi, T., Ceccarelli, A., & Bondavalli, A. (2021). Unsupervised algorithms to detect zero-day attacks: Strategy and application. IEEE Access, 9, 90603–90615. https://doi.org/10.1109/access.2021.3090957

UNDER PEER REVIEW IJUAR