WHITEPAPER

DRAGOS
SAFEGUARDING CIVILIZATION

ONR Office for Nuclear Regulation

Threats & Risks to the United Kingdom's Civil Nuclear Sector

Short-Term, Mid-Term, and Long-Term Guidance for Dutyholders

DR. JACOB BENJAMIN | CONSULTING DIRECTOR

JAN HOFF | PRINCIPAL INDUSTRIAL INCIDENT RESPONDER

FEBRUARY 2025

Executive Summary

This report provides dutyholders in the United Kingdom (UK) Civil Nuclear sector with intelligence to support forward planning and the strategic management of cyber security and information assurance (CS&IA) risks. Its objective is to serve as a catalyst for organisational leadership, encouraging discussions on preparedness to identify and address emerging CS&IA risks effectively. At the request of the Office for Nuclear Regulation (ONR), Dragos examined emerging threats and risks to the UK civil nuclear sector over a three-year assessment and beyond.

This report characterises each threat or risk by its ability to significantly impact industrial operations. Depending on the dutyholder and the use of technology, some risks might be more relevant than others. The adoption of emerging technologies like Artificial Intelligence (AI) and cloud introduces risks that are solely associated with the use of those technologies. Other risks will emerge from the developments in computational power or societal and technological changes.

Effective leadership, independent assurance, and proportionate regulation are key to managing the sector's current and emerging risks. Preparation and mitigation of these threats and risks require augmenting the current defence approach with an additional focus on active capabilities, such as detection, response, and recovery. Dutyholders should participate in a community visibility and collective defence programme centred on the capability of sharing Operational Technology (OT) threat intelligence through recognised partnerships such as the National Cyber Security Centres (NCSC) Cyber Security Information Sharing Partnership (CISP). Information sharing and collective defence allows dutyholders to scale their ability to conduct robust monitoring, supply chain audits, dynamic risk assessments, and cyber incident response exercises. Together, these activities will strengthen organisational resilience and safeguard against evolving threats and risks facing the civil nuclear sector.

Dutyholders should note that some risks expected in the coming years require long-term planning and countermeasures. Among those are risks related to encryption mechanisms and the aging workforce and skill gap. The UK civil nuclear industry should start activities now to be prepared in the years to come.

Dutyholders should note that some risks expected in the coming years require long-term planning and countermeasures.

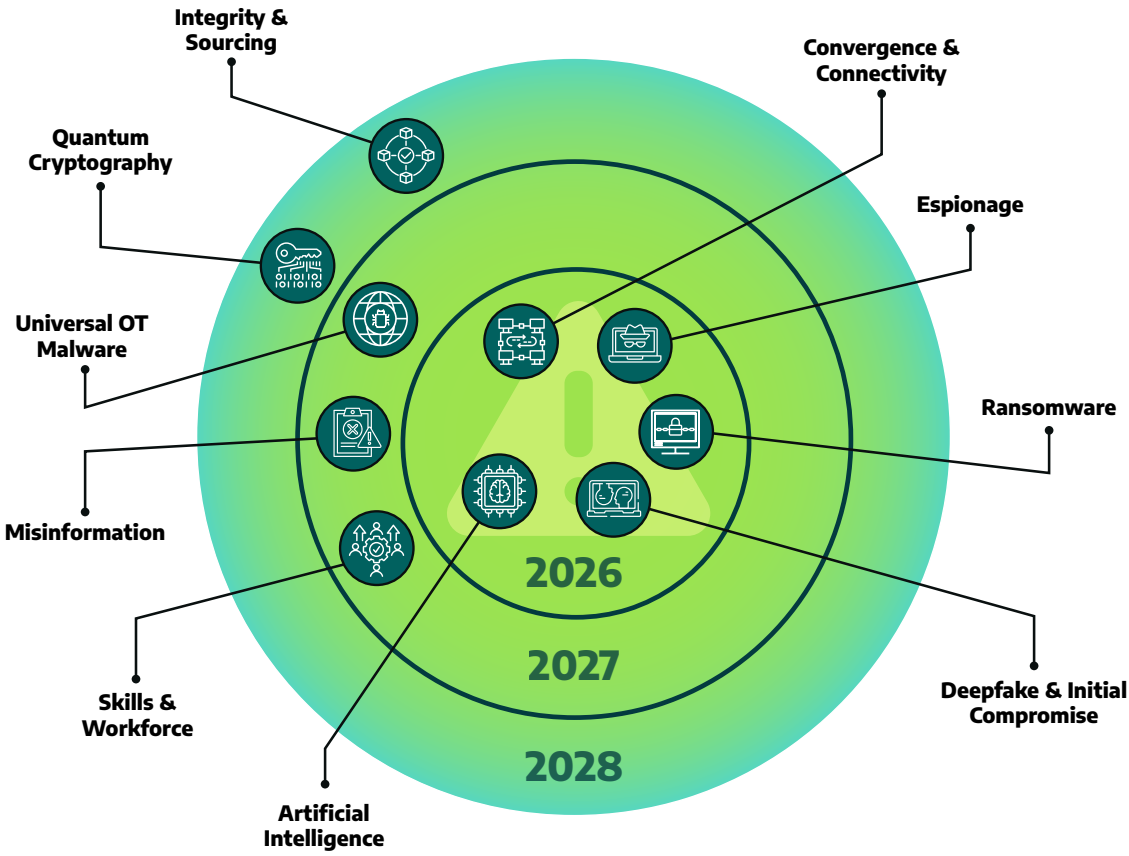

Figure 1: Emerging Threats and Risks

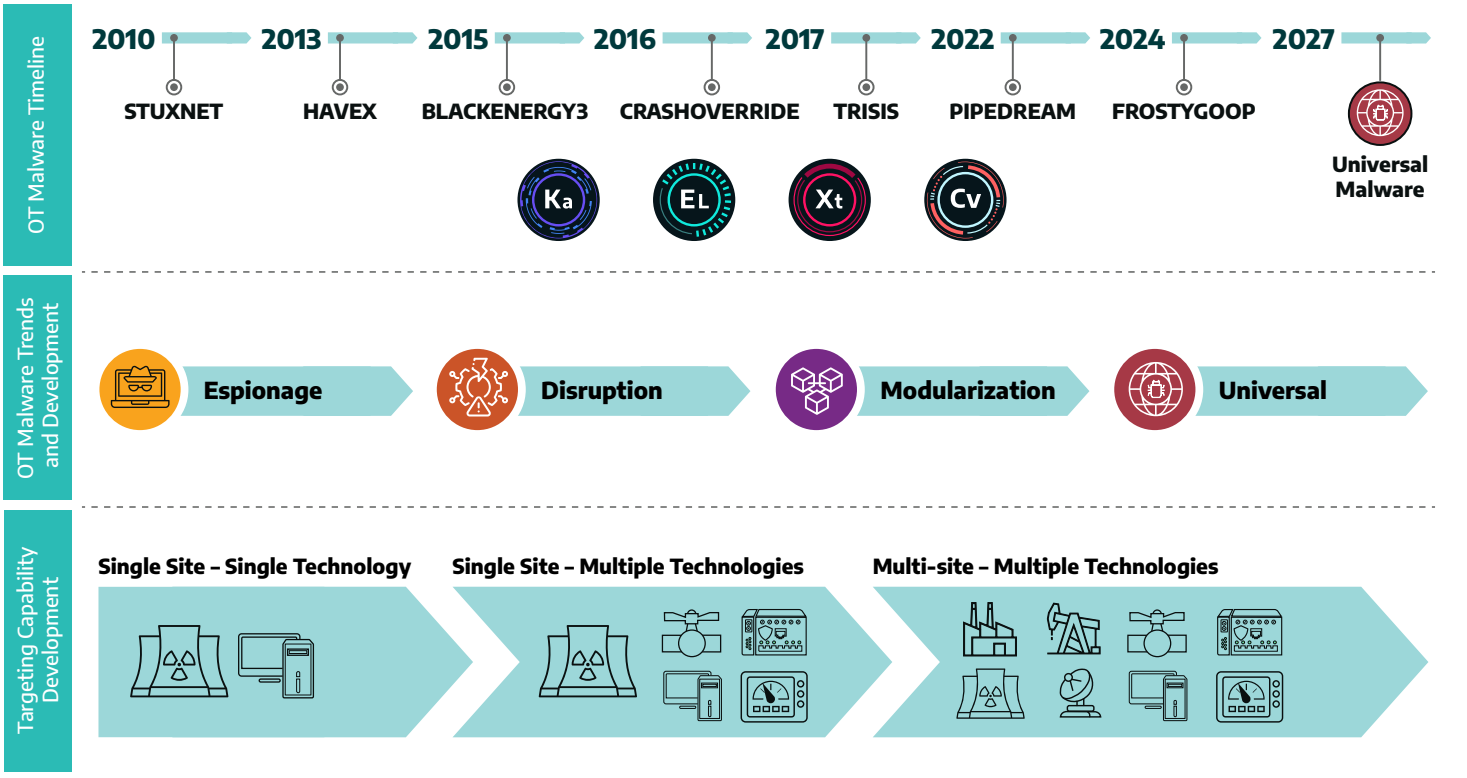

Figure 2: Evolution of OT Malware and OT Cyberattacks

An Emerging Threat Landscape

In recent years, headlines worldwide have highlighted that adversaries targeting operational and industrial technology can disrupt critical infrastructure. Geopolitical events such as Russia's invasion of Ukraine, conflicts in the Middle East, and growing tensions between China and the West have also increased the cyber threat to the civil nuclear sector. Historically, the civil nuclear sector has been risk-averse and conservative in adopting new technology. The concept of "air gap" as a primary security control becomes more and more difficult to retain. As legacy technology receives upgrades or is retired from decades of operating life, the civil nuclear industry will also see a shift in technology and architecture.

Emerging cyber security threats and risks will evolve over the next three years as adversaries leverage technological advances to enhance existing attack strategies and techniques, accelerate the development of new capabilities, and lower the barriers of entry for disruption of OT operations or other high-consequence events.

The emerging threat timeline has three sections: short-term (2026), mid-term (2027), and long-term (2028). A fourth section covers risks associated primarily with Artificial Intelligence (AI) adoption or those that might manifest in 2029 and later. Dutyholders should become familiar with the risks affecting the sector and evaluate the impact and potential impacts.

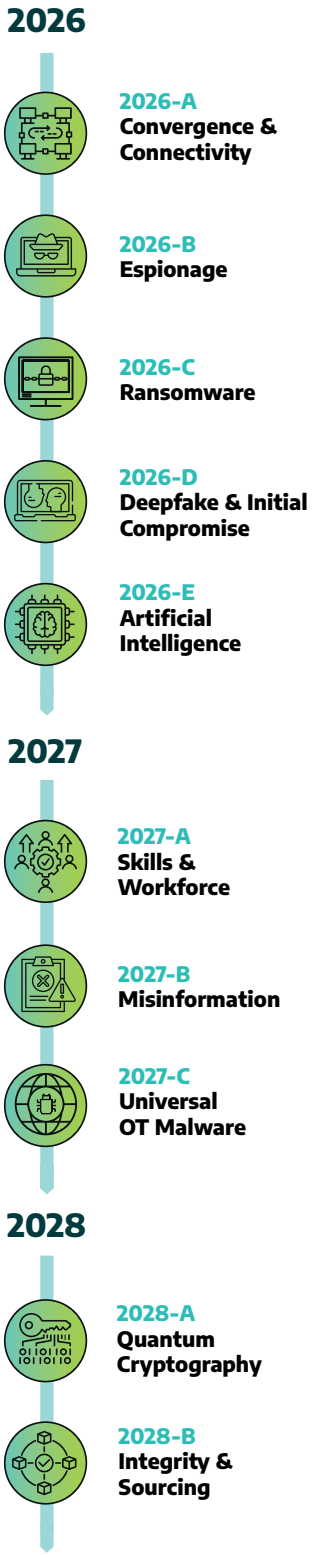

Figure 3: Emerging Cyber Security Threats

1 NCSC: UK Energy Sector: The Cyber Threat Landscape

Short-Term Threats & Risks: Expected By 2026

2026-A: Convergence & Connectivity

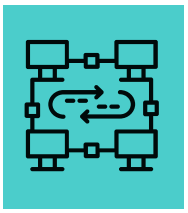

Further technology convergence and hyperconnectivity result in larger attack surfaces and scalable attacks. Conventional hard and software replace custom implementations, requiring an adaptation of security controls.

A shift in technology and architecture for cyber-physical and information-processing systems is underway and continuing rapidly. This includes control systems operating nuclear processes and systems used for planning and business processes. Dutyholders can expect business systems to change faster than operational systems, however they follow the same trend: higher integration, higher standardisation, and off-the-shelf components. Additionally, dynamic computing environments lead to cloud and on-demand computing adoption.

With technology moving to integrated single-solution architectures, dutyholders should adjust to those architectures or invest in maintaining their current designs. Maintaining legacy architectures can be error-prone and expensive. With the workforce change (see risk **2027-A: Skills and Workforce**), dutyholders should plan accordingly.
Vendor acquisitions, technology standardisation, and the convergence of common operating platforms result in control systems and infrastructure for widely different processes becoming increasingly similar. This shift from heterogeneous control systems to more homogeneous architectures enables adversaries to scale their attacks in ways not previously observed. Where in the past, each control system and business system had individual components, which are now interchangeable and differ only in configuration or parameters. Adversaries can codify knowledge and techniques learned previously and directly apply them to new OT attack campaigns with little or minimal adjustments regardless of industry vertical. This convergence is not a new development but the logical consequence of optimisation and synergies in development and operation. Personnel just need to know one type of system – but adversaries also only have to know (and exploit) one type of system.

In line with more homogenous systems, there is a need for information exchange and connectivity between those components within and outside of the organisation's boundaries. Creating a hyperconnected ecosystem, leveraging cloud technologies, Internet-of-Things (IoT), virtualisation, or containerisation, has revolutionised IT by offering scalable resources accessible from anywhere. A rapid deployment of applications and services provides real-time insights and processing of vast data collections. The compounding advantages of these technologies will increase the demand for the civil nuclear sector to adopt more of them.

Greater connectivity (hyperconnectivity) means more cyber-attack entry points, supply chain breach opportunities, and easier malware propagation or lateral movement. Concepts like zero-trust for network access or big data storage databases make conventional security approaches and controls increasingly difficult. The greatest challenge dutyholders will encounter lies in the integration of legacy components (e.g., mainframes) with new technologies (e.g., cloud).

Many modern security solutions rely on cloud-based infrastructures for endpoint protection, threat intelligence, and detection. Adopting operational systems to these IT-centric solutions without disrupting operations or introducing latency is a significant challenge. OT systems, especially within the civil nuclear sector, have been designed for isolation and minimal change. Cloud connectivity contrasts with the traditional OT approach and regulation, necessitating alternate solutions such as application allow-listing or on-premises security deployments. Operating two cyber security stacks (one modern IT stack and a legacy OT stack) is expensive and may lead to uncontrolled security risks.

This risk is inherent to technological change, and dutyholders should establish a strategy for securely connecting previously isolated environments while maintaining the safety and regulatory requirements placed on the civil nuclear industry. Organisations without a proper strategy will lose control of their most important assets and may expose them to unmanaged risks. When legacy environments connect to modern systems, integrity, availability, and confidentiality should not be compromised.

A second approach to counter this emerging risk and manage hyperconnectivity is not to focus on reversing these trends but to embrace the advantages.

With highly connected systems, organisations also gain insights into their security and the potential to identify anomalies. Dutyholders should leverage a community approach to defence. Sharing information on potential threats and best practices in trusted communities can help organisations better assess risks and countermeasures. Recognised partnerships already exist (e.g., CISP, NCSC i100) where vetted organisations can exchange information on anomalies and attacks.

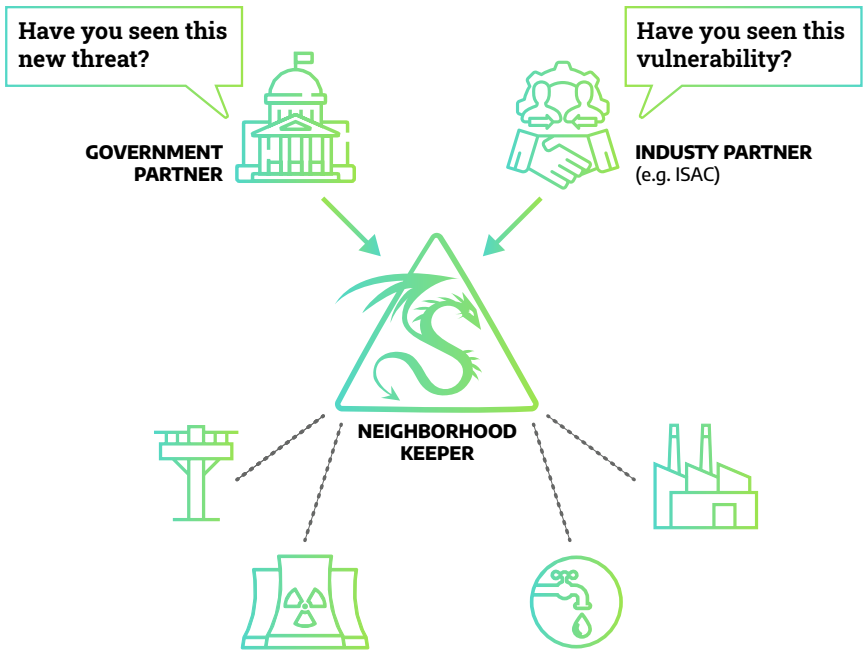

Figure 4: Example of a Community Visibility and Defence Programme

2026-B: Espionage

Intellectual property theft endangers the confidentiality of sensitive nuclear information and industry secrets. Adversaries use reconnaissance for better initial access vectors and have shown interest in intellectual property in the civil nuclear sector. Data exfiltration, as part of or as the goal of intrusions, will increase with higher connectivity and digitalisation.

In contrast to conventional theft, data theft is difficult to detect since the victim does not lose access – the adversary creates a copy of stolen data. Improperly managing the risk from connectivity (see **2026-A: Convergence and Connectivity**) enables adversaries to breach organisations and exfiltrate data. Additionally, the increased amount of data stored and processed makes classifying and protecting that data and their copies more difficult. Using AI (see **2026-E: Artificial Intelligence**) and third-party components (see **2028-B: Integrity and Sourcing**) potentially shifts data processing and storage to less controlled entities. When the stringent controls on data cannot be enforced any more, that data might be an easier target for adversaries. Hack-and-leak operations have also aided in misinformation campaigns (see **2027-B: Misinformation**) and are related to this risk.

The National Cyber Security Centre (NCSC) assesses that the UK civil nuclear sector is an increasingly attractive target for intellectual property (IP) theft and other commercially sensitive information through cyber espionage campaigns. The UK has committed to nuclear research and development and has sensitive nuclear information, making civil nuclear a target.[2] State-sponsored adversaries are often interested in exfiltrating advanced leading-edge research, such as the work on Advanced and Small Modular Reactors (SMR), and fusion energy development in the UK and Europe.[3] Civil technology can also be relevant from a military perspective – examples are nuclear enrichment and fuel repurposing. Therefore, it is highly likely that IP theft through cyber espionage campaigns will remain a predominant cyber threat to the sector.

In the past, state actors have successfully breached civil nuclear institutions globally, and the goals and effects of these intrusions remain unclear. Long-term espionage operations can provide information that may be used for economic and political purposes. The civil nuclear sector is a strategic target for state-sponsored research, reconnaissance, and pre-positioning. Threat groups, such as VOLTZITE[4] and XENOTIME[5], have compromised their targets and focused on data gathering once they established their foothold. Those threat groups also have shown the capability and intent to exfiltrate sensitive information from victim's networks.

Operational data exfiltrated from OT networks may provide an adversary with crucial intelligence to aid in follow-up offensive tool development or attacks against OT networks. Previous activity from XENOTIME has shown that disruptive impacts require extensive research. State-sponsored threat groups have been known to exploit supply chains to gain uncontested access and ample time to plan and develop a multi-stage attack.[6] Supply-chain attacks have affected the civil nuclear industry in the past and will continue to provide a vector into otherwise heavily secured environments (see **2028-B: Integrity and Sourcing**). Espionage and reconnaissance provide the basis for such vectors.

[2]NCSC: UK Energy Sector: The Cyber Threat Landscape; [3] UK Nuclear Fuel Investments; [4] VOLTZITE; [5] XENOTIME; [6] Supply Chain Attacks in OT

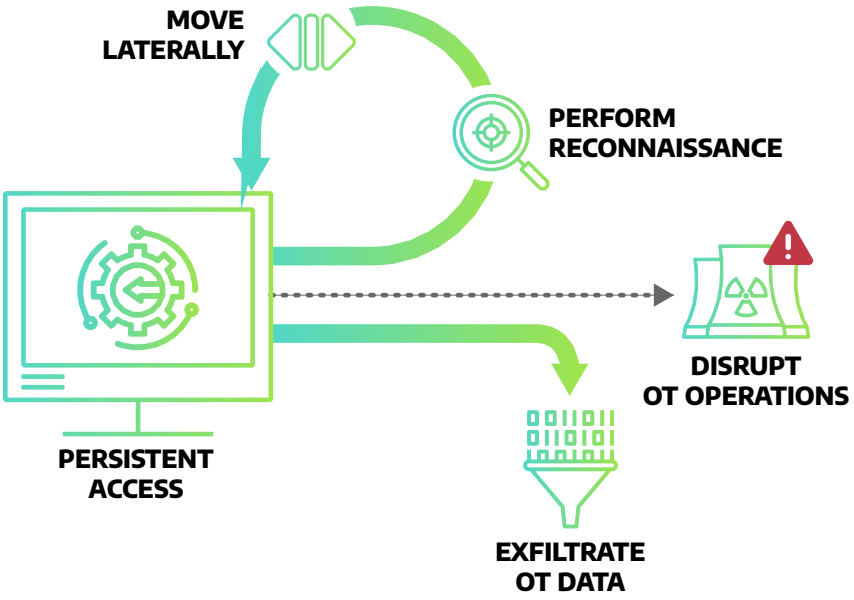

Figure 5: Persistent Access Leading to (State-sponsored) Espionage

In addition to espionage via digital means (i.e., hacking and computer intrusions), economic dependencies can also impact the protection of sensitive information. State actors and governments have invested strategically in the civil nuclear industry. Strategic investments also occurred in suppliers and vendors of control systems and integrators. When third parties establish a controlling stake in a company or strategically influence commercial decisions, protecting intellectual property becomes difficult. Assessing the long-term risks from sucheconomic dependencies is challenging beyond cyber security implications.

Developments in technology, especially the ubiquitous internet access for devices, add to the risk of espionage. Bringing smart (potentially compromised) devices into organisations has never been easier. Among those devices are smart watches, smart glasses, and other Internet of Things (IoT) technology that could collect data, send data, or provide initial footholds into adversaries' target environments. Even though not directly associated with espionage, undisclosed maintenance connections from vendors and radio modems in components could also achieve a similar result. This risk correlates with potential backdoors in hardware and firmware in **2028-B: Integrity and Sourcing**.

Organisations in the civil nuclear sector must first understand their exposure regarding sensitive data, especially for SNI-classified data. Convergence and an increased need to share data makes identifying and classifying data a challenge. Adversaries with access, either digitally or in person, will attempt to gather and exfiltrate data. An appropriate policy framework and an effective enforcing of data handling provides the first step in countering espionage.

Existing practices, including screening third parties for prohibited devices and raising awareness with employees, must accommodate technology advances. Dutyholders should consider the implications when restricting devices and technology. Modern working, especially in the future, might need smart glasses and mobile technology – for augmented maintenance tasks, real-time data entry, and optimised workflows. Conventional data leak protection and data classification might become obsolete when smart glasses and internet-of-things components, compromised or abused, allow for accessing sensitive data.

Another approach to countering espionage and data leaks is increasing the visibility and monitoring of sensitive and OT networks. Identifying and protecting vital information is part of this approach. While monitoring does not prevent espionage, it does allow to identify anomalies and potentially the extent of data that was misused or in access by an adversary. Dutyholders process and store vital information likely to be targeted for theft or exfiltration, such as sensitive nuclear information, research, and operational data. Tagging and flagging transfer of that data can aid in detection and response to aforementioned anomalies.

Stringent access controls and data encryption should accompany all data classification procedures. This includes protecting backups or drafts that would allow an adversary to reconstruct data. Once adversaries have exfiltrated data, they may try to analyse and replicate environments and technologies – either for vulnerability research or reverse engineering.

Using native tools and fileless attacks often evade traditional antivirus software. Living-off-the-land techniques and credential reuse enables these groups to avoid detection. Combining this approach with a slow and steady reconnaissance extends dwell time within a network and anomaly detection. This makes alarms less likely for defenders. Dragos recommends monitoring cross-zone communications between IT and OT networks and utilising behavioural detections engineered to identify the latest applicable tactics, techniques, and procedures (TTPs) and indicators of compromise (IOCs).

2026-C: Ransomware

Ransomware continues to threaten organisations in the civil nuclear sector. Where commodity ransomware can primarily impact less mature organisations, targeted ransomware operators gaining access to dutyholders' networks will employ advanced intrusion techniques to achieve their objectives. New technologies and research fuel the evolution of malware and successful initial intrusions, making successful ransomware operations more likely.

Currently, ransomware represents the most likely disruptive threat to the UK civil nuclear sector and its supply chain, and it is expected to remain so for the next two years.[7] The emerging threat regarding ransomware is the evolution of Ransomware-as-a-Service (RaaS) as cyber criminals integrate AI-powered features and reduce detectable traffic patterns. RaaS has become the dominant operating business model, with almost all the top ransomware strains operating this way. It has lowered the barrier to entry for less skilled cyber criminals, and it has proven difficult for legislators to disrupt RaaS operators and their affiliates. As detailed in the Dragos Year in Review and the NCSC's "The Threat from Commercial Cyber Proliferation", there is a growing convergence of interests between sophisticated adversaries and cybercriminal including hacktivists and ransomware groups.

Advances in artificial intelligence can amplify RaaS capabilities, making attacks more targeted, efficient, and harder to detect. AI may enhance targeting by analysing vast datasets such as public records, social media, and data dumps from compromises to profile organisations and even predict their ability or willingness to pay ransom. AI chatbots increase the efficiency of attacks by automating negotiations or assisting in triple extortion strategies with tailored messaging to victims. Generative AI can help ransomware adapt and modify its payloads over time, increasing its ability to evade detection and maintain persistence.[8] Modern malware, including ransomware, often does not "touch disk" anymore but only resides in memory. Conventional forensics and detection tools have to constantly scan memory to identify potentially code – leading to additional computing overhead on already constrained industrial environments.

Ransomware prevention and response mitigations for industrial environments includes a mix of people, processes, and technology. Dutyholders should maintain regular, encrypted and offline backups of critical data, with periodic restoration testing. Operational uptime is paramount in the civil nuclear sector and having access to recent, reliable backups is key to a swift restoration of operations. With or without AI features, ransomware groups will continue using living-off-the-land techniques to escalate privileges, gain persistence, and establish command-and-control (C2) channels. Asset visibility and network monitoring are essential for detecting ransomware early in its lifecycle – such as the initial C2 beacon or before data is fully exfiltrated from the network.

Remote access should require multi-factor authentication (MFA). Sessions and activities have recording and controlled flow through a hardened bastion host with increased logging and monitoring applied.

Conducting tabletop exercises with a ransomware scenario allows defenders to practice containment, eradication, and recovery processes, walkthrough the incident response plan (IRP) and validate roles and responsibilities in a low-stress environment.

[7]NCSC: UK Energy Sector: The Cyber Threat Landscape; [8]AI-Powered Cyberattacks

2026-D: Deepfake And Initial Compromise

Evolving deep-fake, high-interaction AI bots and augmented phishing campaigns allow adversaries to be more effective at initial intrusions and automate the circumvention of conventional security controls. A further surge in successful initial access campaigns requires organisations to establish defence-in-depth and prepare response actions.

Advances in AI-powered chatbots and deepfakes increase the effectiveness of common initial access techniques such as phishing and social engineering against humans. AI-driven chatbots can analyse publicly available information, such as social media or company websites, and craft highly personalised messages to target individuals or organisations. Recent experiments have revealed that a significantly higher percentage of users clicked on links in AI-generated phishing emails than human-written ones.[9] Adversaries are beginning to use AI tools, such as deep learning for impersonation. Deep learning uses neural networks to create convincing images, audio, and video of trusted personnel, known as deepfakes. Deepfakes produce believable outputs beyond simple face swapping or photo application filters. Cybercriminals have successfully utilised deepfake technology to transform publicly accessible videos and other media into realistic replicas of video conference call attendees, resulting in significant financial loss.[10]

Deepfakes could be weaponised for more sophisticated impersonation attacks, potentially leading to insider threat incidents where attackers gain unauthorised access or influence high-level decisions within organisations by impersonating executives or employees. In a long-term plot, adversaries may use deepfakes to circumvent personnel vetting. Video and audio deepfakes could support fabricated educational, professional, or personal backgrounds, making the individual appear highly qualified and trustworthy. Deepfake-generated photos and videos can build a compelling online profile, bolstering credibility during background checks. A live deepfake video call could enable the adversary to pass interviews and remote onboarding procedures. Once inside, deepfakes can help the insider solidify their position and access by falsifying authority or forging communications to enable the creation of false identities or credentials, cover up their activities, or redirect susception for security breaches or policy violations. Adversaries are currently limited to templates featuring human avatars reciting voice-to-text speech and face swap tools on existing videos. Still, they are investing in generative AI, and dutyholders should prepare for that evolution.[11]

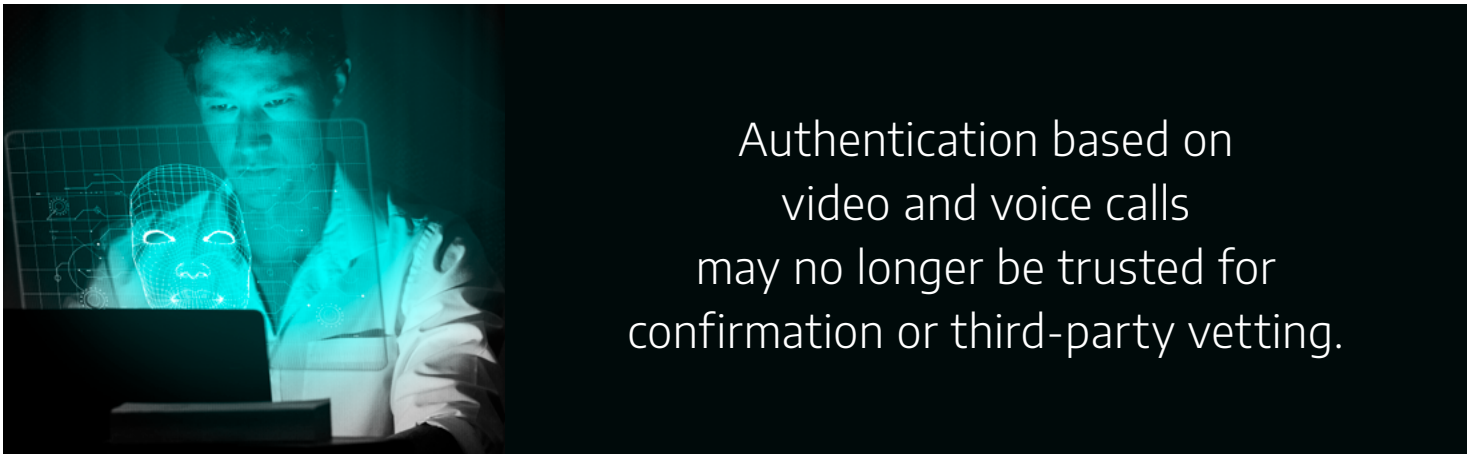

[9]Applying AI to Targeted Phishing Attacks; [10] Deepfake Scam; [11]Threat Actor Generative AI Capabilities

Initial access is the prerequisite to further stages of an attack and is typical of all cyber attacks regardless of the adversary's goal.[12] It consists of various techniques to gain an initial foothold within a network as they strive to escalate privileges, expand access, and gain persistence. While some adversaries see initial access as the means to perform further actions on objectives, such as espionage or ransomware, several OT threat groups focus on developing and obtaining initial access against industrial organisations.[13] For example, KAMACITE[14] typically provides access to other entities, such as ELECTRUM[15] to execute OT disruptive effects while also maintaining long-term access.

Data exchange mechanisms and references to digital media will continue to change. With more information digitalised and personnel needing quick and mobile links will continue to develop. Among those are near-field communication (NFC) tags or QR codes (Quick Response): barcodes or wireless transponders that encode information such as uniform resource locator (URLs), text, or commands. They are widely used for convenience and ease of scanning using smartphones and other devices. However, their simplicity also makes them a vector for security risks when exploited by adversaries. The primary concern with QR codes is counterfeit codes in phishing attempts, furthering social engineering efforts, or distributing initial access malware. QR codes are often used in multi-factor authentication (MFA) enrolment, maintenance manuals, components tags, or procedures within OT. A tampered QR code in a manual or procedure could lead to malicious configurations that disable or enable the settings of critical equipment. In a less likely but still plausible scenario, emergency response procedures or manuals could be manipulated to interfere with crisis management and recovery operations.

Given that initial access is a prerequisite to any cyber attack, eliminating or reducing initial access vectors can successfully thwart many intrusion attempts. Standard initial access mitigations such as network intrusion prevention systems (NIPS), email filtering, endpoint detection and response (EDR), and awareness training will likely remain effective against these enhancements.

QR codes and NFC tags do not present an immediate threat. However, operators need to ensure the integrity of such technology. Where barcodes could only store limited data, QR codes and tags are more versatile. Dutyholders should continue to digitalise their processes but ensure that manipulation or unavailability of such technologies does not impact operations. Replacing or overwriting tags requires physical access like insider threats.

The existing insider threat mitigation programme within the civil nuclear sector, with elements like personnel vetting, two-person rule, and security awareness, is effective. Additional measures, such as having more visibility in the OT and IT environments, will enable an organisation to detect unusual activity faster. Example alerts may include creating a baseline and detecting when accounts log in outside of their normal shift hours or when employees badge in outside of normal physical areas, and identifying when employees install new software, run new scripts, or access new network resources. Leveraging job rotations and separations of duties can make it more difficult for a single insider to cause large-scale damage. In cases where this is impossible, additional oversight, such as logging and auditing user interactions, especially with crown jewel systems, is strongly advised.[16]

[12]SANS ICS Cyber Kill Chain; [13]How Dragos Activity Groups Obtain Initial Access into Industrial Environments; [14]KAMACITE; [15]ELECTRUM, [16]Insider Threats in OT

2026-E: Artificial Intelligence

Artificial Intelligence combines elements from machine learning, language models, and universal problem-solving. The more universal this technology becomes, the less deterministic it will be. Individual application's source data requirements might threaten data confidentiality, the lack of reasoning threatens the integrity. Dutyholders need to ensure that technology still fulfils the safety and security requirements.

The term AI has ambiguity. It is not just one technology but a collection of technologies. Machine Learning has direct applicability in today's civil nuclear sector solving problems, designed to excel at specific tasks such as classification or predictive modeling. Generative AI can create new content, including pictures, texts, and potentially training and operations documents relevant to dutyholders.

Many AI models have intermediate layers that process input data before producing an output. These layers are "hidden" because how they transform and extract patterns from data is not directly visible or interpretable to users. Hidden layers are a technical challenge to be solved prior to the adoption of AI for making critical decisions. Dependency on third-party AI models and platforms makes supply chains increasingly attractive targets.

The final stage of a universal AI is too far in the future to influence dutyholders. The lack of transparency and reasoning, in combination with increased data processing associated with generative AI, introduces new risks.

- Organizations planning to use AI should conduct a careful trade-off analysis of the risk and benefits.
- Tailor training programs to all levels of organisations, emphasising AI's impact on nuclear security.
- Enhance trust/confidence in AI decisions by establishing clear chains of responsibility.

The section **Threats and Risks related to AI in 2028 and Beyond** provides further outlook on AI and its implications. Adaption of AI will vary in different industries and the dutyholder's role. The higher the dependency or implementation of such features, the more relevant are protective controls and understanding AI outputs.

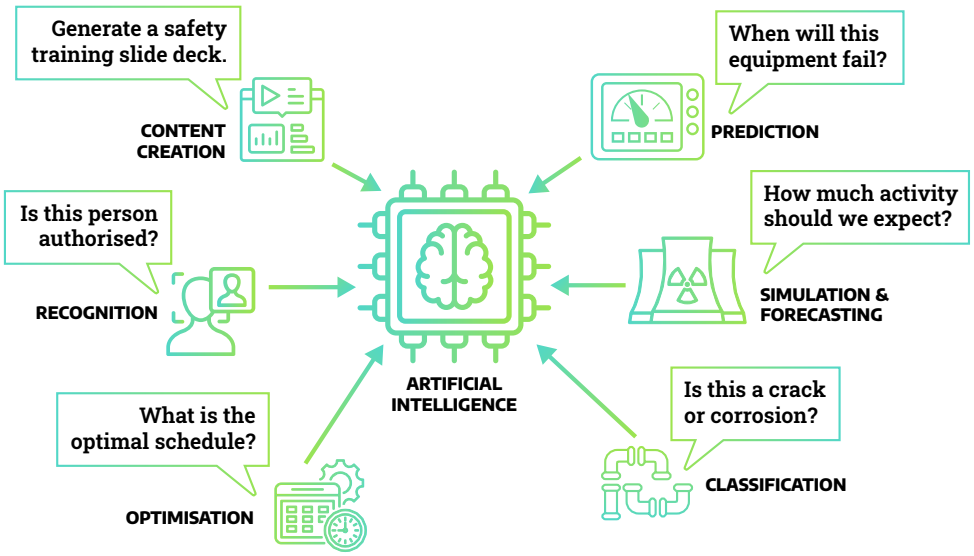

Figure 6: Use Cases for AI and Possible Areas for Abuse/Risks

Mid-Term Threats & Risks: Expected By 2027

2027-A: Skills And Workforce

Cyber security skill shortages and workforce shortages will impact the ability to plan, implement, and maintain cyber security safeguards. A retiring workforce and a new workforce without legacy technology knowledge will result in gaps and potential attack vectors for adversaries. Organisations will see a challenge in securing existing systems and in establishing new ones.

The civil nuclear industry will encounter a shortage of qualified workforce that can implement, operate, and manage systems and components. In the nuclear field, components need to operate for decades. This also means that the workforce needs to be able to understand and manage these components in the coming years. With every year passed, this risk increases to impact operations and to be able to maintain security controls. Considering new systems' convergence, connectivity requirements, and the increasing need to share information, adversaries can exploit non-maintained systems. They will evade detection since personnel for monitoring and response will be scarce. Incidents like log4j, the CrowdStrike bluescreens, or WannaCry have shown that human resources are key in assessing risks and recovering from incidents.

The workforce and skill shortage originates from two factors:

1

Existing workforce leaving the industry due to age or regular turnover.

Many administrators and technical experts involved in the initial commissioning of nuclear systems have reached their pension age, and knowledge is lost. Additionally, companies often do not backfill positions that leave the organisation.

2

New workforce joining the job market without training on legacy systems and the complex environments of the civil nuclear industry.

Universities and conventional education do not cover cyber security curricula bridging legacy and modern systems.

A lack of cyber security experts in the UK's nuclear industry leads to a higher reliance on third parties, consultants, and vendors. Recognised professional bodies and certifications can help establishing a baseline and level of knowledge/skill required in the sector. Vetting third parties is already a challenge today. Relying on third parties due to a lack of workforce can introduce critical dependencies and risks of an uncontrolled insider. When organisations need additional workforce, e.g., to handle incidents or react to changing risk profiles, surge capacities from third parties might be limited.

Introducing AI into workflows and automating tasks will also affect the workforce. Machines will replace repetitive tasks, and personnel can focus better on non-automatable tasks. At the same time the workforce will lose skills and knowledge they have and need when machines encounter challenges. This skill atrophy will gradually occur and is

hard to counter. Ideally, dutyholders can use the positive aspects (automation and efficiency) to partially counter the negative effects (skill atrophy and lack of experience). Identifying and training for core skills and those needed when AI fails is key to handling this aspect of workforce risk.

The civil nuclear industry should start as soon as possible to identify the required skills of their cyber security workforce. Changing curricula and training future employees will take time, and once companies cannot compensate for knowledge gaps and workforce shortages, it is too late. Companies should consider the time it takes to onboard and develop cyber security workforce for the civil nuclear sector. Clear succession planning and workforce management are key to countering the risk of an ageing workforce and legacy components.

Addressing this risk requires collaboration from all sector participants: the UK government, the industry, and academia. When the workforce and technology age, all three parties will see an effect in their area of responsibility.

The approach to counter this risk is two-fold:

1

Develop the existing workforce with the skills to manage and secure new systems and technology.

This ensures that introducing new technology into legacy environments receives appropriate consideration. Experts in legacy technologies need to be upskilled to secure converging technology against new threat vectors and threat actors.

2

While experts for legacy components and nuclear processes are still working in their corresponding roles, they should be encouraged to share and document their knowledge.

Cross-skilling new employees for the intricacies of civil nuclear and the capability to handle legacy systems will counter the risk of losing knowledge when senior cyber security professionals leave the organisations.

Lastly, the civil nuclear industry should be an attractive field of employment. Engineers and cyber security professionals should be incentivised to join the industry, which requires competitive pay, a positive outlook, making an impact, and a good public image. Organisations should adopt a crawl-walk-run approach when developing the workforce of the future.

2027-B: Misinformation

Cyber-enabled misinformation and psychological operations coincide with hybrid warfare already present in today's geopolitical conflicts. Technical developments like AI and hyperconnectivity lead to the effects of scale. The civil nuclear sector is a controversial topic for society and a viable target for influence operations. Dutyholders need to establish means to identify and handle such campaigns. Misinformation will not directly impact operations but has indirect impacts on the sector.

Hacktivists increasingly aspire to disrupt Western OT and draw attention to geopolitical or social causes They propagate fear, uncertainty, and doubt (FUD) to influence perceptions and create a narrative of instability. As detailed in the Dragos Year in Review, there is a growing convergence of interests between sophisticated adversaries and hacktivist personas. Nuclear energy remains a highly controversial topic, often dividing public opinion due to concerns about safety, environmental impact, and its association with catastrophic events such as Chernobyl and Fukushima.

While nuclear advocates emphasise its role as a reliable, low-carbon energy source crucial for combating climate change, opponents focus on risks like radioactive waste, potential accidents, and nuclear proliferation. A well-orchestrated misinformation campaign could devastate the industry in such a polarised environment. False narratives, AI-generated deepfakes, or manipulated data could amplify fears, erode public trust, and fuel anti-nuclear sentiment. This could delay projects, disrupt operations, and undermine confidence in regulatory bodies and operators, hindering progress toward energy goals.

Influence operations and misinformation also directly effects workforce. A negative public perception of the industry hinders workforce development (i.e., less graduates/professionals willing to join civil nuclear) and adversely influences the existing workforce (i.e., professionals leaving for other sectors).

To prepare for this emerging threat, dutyholders should take several proactive measures and communicate with peers and the regulator if they suspect malintent. First, they should establish clear protocols for identifying, assessing, and countering misinformation campaigns. These protocols should define the roles and responsibilities of crisis communication teams, CS&IA leads, and senior executives.

Training employees to recognise false information and promptly report suspected campaigns to the appropriate teams is one approach to counter misinformation. Maintaining transparency is another step; organisations should regularly share credible updates with the public, regulators, and stakeholders about operations, risks, and safety measures. Finally, building relationships with trusted media outlets can provide valuable allies in disseminating accurate information, helping to combat misinformation effectively.

Awareness is the most effective method against misinformation and influence campaigns. The civil nuclear industry needs to monitor the media and communicate proactively with the public, employees, and cross-organisational stakeholders. Sharing information with industry peers and with the government can help counter potential issues and raise visibility.

2027-C: Universal OT Malware

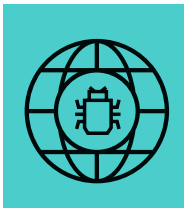

Scalability and repeatability have become themes that will enable new, faster compromises. The trend towards configurable, modular malware frameworks will lead to the emergence of universal OT malware families. Dutyholders will encounter OT-specific malicious code. Adversaries can reduce the time to deploy functioning code impacting industrial environments.

Over the last 15 years, OT cyber attacks have ranged in sophistication, repeatability, and impact. Adversaries have leveraged traditional IT malware and OT malware to achieve these effects, ranging from physical destruction to immediate limited disruption. Analysis of current OT malware reveals that adversaries seek to achieve impacts far greater than immediate disruption by undermining fundamental aspects of process integrity. While such malware was highly specific in the past, adversaries have moved to modular and universal base code that can impact many control components and industries. This advancement in adversary objectives and motives will lead to OT malware families, different versions of the same code, tailored for specific targets or industries. This code base of shared characteristics, functionality, and intent will be like what is already common in traditional IT malware. The slight modifications or configurable aspects will not only decrease the time from initial intrusion to impact but assist adversaries in evading detection.

Adversaries are developing tools that make reusing and deploying against many targets easier. Their OT targeting capabilities have progressed from a narrow focus on a single site or technology to multiple sites and technologies. CRASHOVERRIDE was an initial step of OT malware towards modularity. It included an OT malware base that could load and execute modules based on the target environment's protocols, enabling adversaries to target multiple technologies within a single target environment. PIPEDREAM is fully modular, with the ability to target multiple sites and multiple technologies by creating by abusing well known or common OT protocols like CODESYS, Modbus, or similar. Using well known protocols increases the odds that their tool will be capable in target environments and the modularity eases capability expansion, reuse and deployment. FROSTYGOOP leveraged a configuration file empowering adversaries to reuse the tool with no code changes and the ability to deploy it against any Modbus target.

The timescales needed to conduct a disruptive or destructive OT attack is shortening. Adversaries are exfiltrating OT data and can leverage AI to expedite analysis and targeting. The resulting analysis can be used to craft highly configurated OT malware from the extensible OT malware frameworks. Consequently, the OT industry can expect to observe OT malware with the same base code, but plugins loaded for specific targets. An early example of this that may indicate a trend is with CRASHOVERRIDE and INUSTROYER2. In 2016, the OT malware framework, CRASHOVERRIDE was discovered with a base launcher and modules for IEC-101, IEC-104, OPC-DA, and IEC-61850 protocols. In 2022, INDUSTROYER2, a more targeted variant of CRASHOVERRIDE was found. It had minimal code updates and a narrow focus on just the on the IEC-104 protocol.

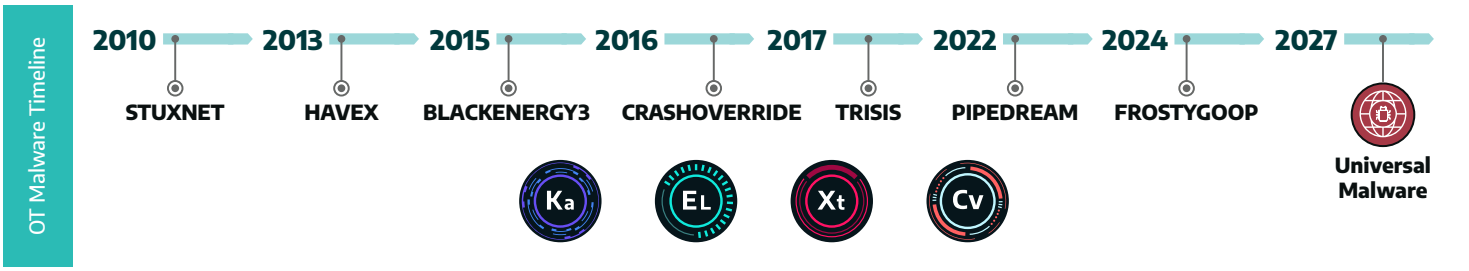

Figure 7: Timeline of OT Malware and Attacks

The most effective protections for this emerging threat are the implementation of visibility and monitoring and a well-practiced OT-specific incident response plan (IRP). It is essential to have procedures for operating with a hampered or degraded control system. Understanding the entrances into these critical network segments, even those employed behind airgaps or one-way deterministic devices is essential to an effective defence-in-depth strategy. Regular supply chain audits and removable media programs will help, but not entirely mitigate the risk of malware bypassing isolation devices.

Monitoring the entrances and information flow chokepoints enables effective response and recovery. Dutyholders should monitor East-West traffic in OT networks with OT protocol-aware technologies, looking for modifications outside of maintenance periods in addition to ingress and egress points of the network.

While OT malware may not require C2 communication, it would likely still interact with the OT assets via native protocols, which may still be detected via anomaly and behaviour-based detections.

Deeply isolated network segments usually have regular traffic patterns and are easily baselined by monitoring solutions with minimal concerns about false positives.

Employing baselines has been shown to be effective against detecting malicious activity for which there are not yet signatures or indicators of compromise. It is highly unlikely nation states currently have the intent to conduct a highly destructive attack against the UK civil nuclear sector, but the timescales needed to conduct a disruptive or destructive attack is shortening and an attack causing temporary and limited disruption is a realistic possibility. Scalability and repeatability have become themes and trends indicate a continued progression towards configurable, modular malware frameworks producing OT malware families and variants.

Long-Term Threats & Risks: Expected By 2028

2028-A: Quantum Cryptography

Quantum-enhanced cryptographic breakthroughs endanger the confidentiality of stored data and established communication encryption methods. Dutyholders need to plan and phase out potentially insecure encryption. Quantum risks are still theoretical today, but once they are proven practical, they will immediately impact various applications and systems.

Researchers and corporations have already developed quantum computing devices. These devices perform computations differently than conventional algorithms making certain tasks significantly faster. Among those tasks is large number factorisation – the foundation of asymmetric cryptography. Such cryptography relies on the complexity of performing such tasks, and organisations should consider algorithms, like Rivest-Shamir-Adleman (RSA), broken once quantum computing has reached sufficient maturity.

Adversaries could use quantum computers to break traditional encryption methods, rendering sensitive data collected today vulnerable – and rendering conventional encryption algorithms ineffective. Even though necessary computing power is not publicly available today, expected computing development and a "collect now, decrypt later"-approach pose a threat to industries with long-living infrastructure. Current estimates for breaking one of the critical algorithms, like RSA, are 10 years from now. Adversaries may collect and store data today and decrypt it at some point in the future. Such attacks are likely to only be worthwhile for very high-value Sensitive Nuclear Information (SNI) or other classified or proprietary information. Theft of this information, even in its encrypted state, is a risk due to ongoing reconnaissance and pre-positioning activity by state-sponsored threat groups such as VOLTZITE and XENOTIME. It is likely that state-sponsored actors will secretly have access to quantum computers of sufficient power before the general public does. This should influence the timeline of dutyholders when planning for this risk.

Migration to post-quantum cryptography (PQC) is a multi-year activity and may take over a decade to complete – especially considering legacy components in the civil nuclear sector. Some components might lack an upgrade capability, and data encrypted with today's algorithms are susceptible to decoding once quantum computing reaches general availability. Identify systems that store SNI or other high-value targets for long-term encryption breaking. PQC upgrades can be planned as systems are being updated or replaced. NIST has published three PQC algorithms (ML-KEM, ML-DSA, SLH-DSA). ML-KEM (Kyber) and ML-DSA (Dilithium) are algorithms standardised by the National Institute of Standards Technology (NIST) that are suitable for general-purpose use.

Quantum computers do not significantly impact the security of symmetric cryptography, and existing symmetric algorithms with appropriate key sizes can continue to be used. Quantum computers do not optimise problems for common symmetric algorithms and as such will not directly require actions. Organisations in the civil nuclear industry should continuously assess their use of cryptography and consider future developments in encryption and decryption capabilities. Dutyholders have to perform a continuous evaluation of key lengths and algorithms, since computing power will increase – independent of quantum computing.

2028-B: Integrity & Sourcing

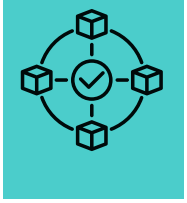

Using off-the-shelf products and commodity components as part of convergence will result in additional risks. Hardware and software sourced from third parties will be harder to vet or trace to their origin. Organisations will not be able to detect component modifications, potentially leading to unexpected behaviour or security flaws (e.g., backdoors). Using off-the-shelf products will allow attacks at scale and bypass conventional integrity vetting processes.

Infrastructure and systems in the civil nuclear industry are integrated. Full traceability of components and software sources will become more difficult with higher integration and using off-the-shelf components. Network devices and services already contain more functionality and code than can ever receive validation. Changes in hardware design and printed circuit board (PCB) layouts are difficult to detect. The same applies to proprietary firmware utilising third-party libraries. Manufacturers have to either flash their own firmware after production or rely on a secure manufacturing process, where devices receive unmodified firmware. In addition to manipulation during manufacturing, code validation during development is challenging.

Hardware components are often manufactured in Asia (primarily China) as part of contracting with semiconductor fabrication and assembly companies. While hardware vendors retain intellectual property, they have limited control over the manufacturing process. With technology convergence and the use of off-the-shelf components, organisations will lose the traceability of integrity and origin of those components.

Hardware-based supply chain attacks involve compromising components during the manufacturing or distribution process. This could be the insertion of malicious chips, altered firmware, or backdoors that lead to logical flaws or attack vectors. Although “hardware implants” have mainly been exaggerated or untrue in the past, they remain a viable attack vector for well-sourced adversaries.

In addition to tampering with hardware, firmware remains a vital component of nearly every device – especially in the OT. With increasing complexity and third-party components in firmware, flaws become more likely. If an adversary gains control over third-party libraries or the development process, malicious firmware could make its way into devices used in the civil nuclear industry.

Another issue introduced by convergence and off-the-shelf products is counterfeiting. Even if manipulation does not occur maliciously, operators might implement components that are not as they expect them to be. Initially, products might work as expected, but with longer operating time performance degrades, unexpected errors occur, or updates are incompatible.

Only large organisations have resources for validating the integrity of hardware and firmware. Even then, not every component can receive checks and security testing. If adversaries added additional chips to a PCB, those might be visible during an X-ray, but if chips are replaced, it is near impossible to detect. Government-sponsored entities and agencies have been using manipulated hardware to perform espionage since the Cold War. Adversaries will likely utilise supply chain weaknesses and attack vectors to bypass conventional security controls. In a globalised world and outsourced manufacturing, maintaining control over the lowest levels of hardware and software will be a challenge that dutyholders have to address.

Dutyholders should establish clear protocols for supply chain validation and require a bill of materials (hardware and software) as part of their sourcing. Audits and individual assessments allow the detection of anomalies. This must be a sector-wide approach and cannot rest on individual organisations. The United States has established regulations banning certain Chinese manufacturers from being used in critical infrastructure. Even if today's products are not inherently malicious, future updates and modifications could subvert security controls in place today.

The existing supply chain programme within the civil nuclear sector, with elements like audits, inspections, vetting of suppliers, and collaborative efforts among regulators, operators, and suppliers, is effective. Additional measures, such as having more visibility in the OT environments, will enable an organisation to detect unusual activity faster. Organisations should ensure that they can security test devices and that components have not been tampered with.

Threats & Risks Related To AI In 2028 & Beyond

Many vendors for systems and equipment used in the civil nuclear sector are considering incorporating artificial intelligence (AI)/machine learning (ML) features into their products striving for improvements in efficiency, safety and operational effectiveness. Several system functions such as visual inspection and quality control, surveillance and security, energy management systems, hazard detection, or predictive maintenance can leverage AI/ML models for classification and analysis of images, audio, video, sensors, etc. It is unlikely that these features will be widely adopted in the civil nuclear sector before 2030. However, if they are eventually adopted new risks such as adversarial AI attacks, data poisoning, and hidden layers are likely to emerge. Dutyholders in research and development are more likely to implement AI features than operators in production.

Adversarial AI is where AI algorithms are manipulated into making incorrect decisions on classifications or analysis. This deception can mislead operators into making incorrect decisions, such as overloading equipment or neglecting necessary maintenance measures, potentially causing damage to crown jewel assets.

Dependency on third-party AI models and platforms makes supply chains increasingly attractive targets. The most likely attack vector for supply chain attacks is compromised updates. Regular updates to AI software and models are a critical part of maintaining performance and addressing vulnerabilities. Attackers may attempt to compromise these updates to introduce malicious code or vulnerabilities. Additional attack vectors such as embedding backdoors in the AI models and tampering with data pipelines. According to a 2023 Gartner report, nearly 30 percent of AI-enabled organisations, experienced data poisoning attacks.[17] This risk from model poisoning is that adversaries embed specific triggers in training data, causing the models to behave incorrectly when those triggers appear in real world inputs. This could potentially lead to unnecessary operation shutdowns or the failure to identify real issues, potentially those with significant safety concerns.

Many AI/ML models have intermediate layers within a neural network that process input data before producing an output. These layers are "hidden" because how they transform and extract patterns from data is not directly visible or interpretable to users. In addition to a lack of transparency, training data limitations and overfitting are significant concerns when working with hidden layers in AI. If the data used to train the model contains biases, then those biases will be embedded into the model's decision-making process. For example, if a model is trained on biased historical data that is skewed towards underreporting of minor incidents or overlooking specific components, the

[17]Gartner

AI might prioritise certain types of equipment or safety protocols over others. This could result in the AI/ML model failing to identify emerging risks in underrepresented areas. The more hidden layers an AI model has, the more opportunities there are for the model to memorise the training data rather than generalise, causing the model to overfit and lose robustness when faced with new data.

While adopting AI/ML features may lead to the theoretical attack vector presenting risks, there have been no documented instances of such attacks successfully compromising operationally deployed equipment or causing widespread damage. Based on publicly available information, adversarial AI attacks targeting equipment used in the civil nuclear sector can currently be classified as non-events. For the civil nuclear sector, data model poisoning and hidden layers in AI may be more appropriately categorised as a technical challenge to be solved prior to its adoption for use related to making any critical decisions, rather than an emerging risk.

References & Appendix

Office For Nuclear Regulation

The Office for Nuclear Regulation is the UK's independent nuclear regulator and competent authority for nuclear security. ONR regulates security arrangements to ensure that the civil nuclear sector adequately protects sensitive nuclear information and industrial systems against cyber-attack and remains resilient to emerging threats as part of their broader mission to protect society by securing safe nuclear operations.

This paper has been created and published by ONR working in partnership with Dragos and is provided for guidance and information purposes. It does not constitute official ONR guidance or regulatory requirements. Further information can be found at **www.onr.org.uk**

View Now!

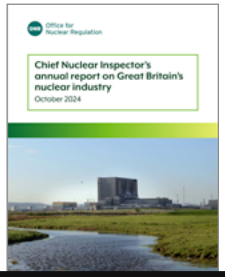

View Now!

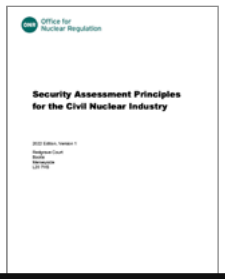

View Now!

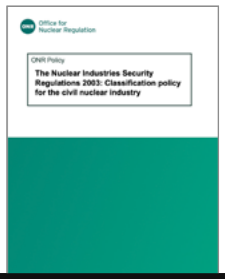

View Now!

View Now!

View Now!

Published by Dragos:

View Now!

View Now!

View Now!

View Now!

View Now!

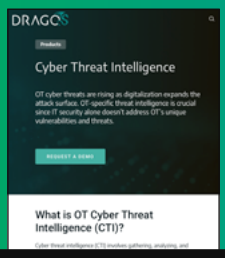

View Now!

ABOUT DRAGOS, INC.

Dragos, Inc. has a global mission to safeguard civilization from those trying to disrupt the industrial infrastructure we depend on every day. Dragos is privately held and headquartered in the Washington, DC area with regional presence around the world, including Canada, Australia, New Zealand, Europe, and the Middle East.

Learn more about our technology, services, and threat intelligence offerings:

Request a Demo | Contact Us

ONR Office for Nuclear Regulation

ABOUT OFFICE FOR NUCLEAR REGULATION

ONR are the UK's independent nuclear regulator and competent authority for nuclear security. ONR's mission is to protect society by securing safe nuclear operations.