
Professional Certificate in Operational Technology Engineer (United Kingdom)

Risk Management in Operational Technology

Risk Management in the context of Operational Technology (OT) is the systematic process of identifying, analysing, evaluating and treating risks that could affect the availability, integrity and confidentiality of industrial control systems. It differs from traditional IT risk management because OT environments are tightly coupled with physical processes, safety functions and often have real-time constraints. The language used in this discipline includes many specialised terms that learners must master to communicate effectively with engineers, managers and regulators. The following exposition defines the most important vocabulary, illustrates each concept with practical examples, and highlights typical challenges that arise when applying the terms in a live OT setting.

An asset is any component that has value to the organisation. In OT, assets include physical devices such as programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA) servers, distributed control systems (DCS), human-machine interfaces (HMIs), and the sensors that feed them data. An asset may also be intangible, for example the process knowledge held by an operator or the software configuration that governs a safety instrumented system (SIS). Understanding the full asset inventory is the first step in any risk assessment because unknown assets cannot be protected.

A vulnerability is a weakness in an asset that could be exploited by a threat. Vulnerabilities in OT often stem from legacy equipment that cannot be patched, insecure communication protocols such as Modbus or DNP3, or mis-configurations introduced during commissioning. For instance, a PLC that runs an outdated firmware version may contain a buffer overflow defect that allows an attacker to execute arbitrary code. In many cases the vulnerability is not a technical flaw alone; it can also be a procedural lapse such as the absence of a documented change-management process.

A threat is any potential cause of an unwanted incident. Threats in OT can be external, such as a nation-state actor targeting critical infrastructure, or internal, such as a disgruntled employee who knows the network topology. Threats are characterised by their source, motive and capability. An example of an external threat is a ransomware group that gains access through a phishing email, escalates privileges, and then encrypts the data historian, halting production. An internal threat might involve a maintenance contractor who inadvertently introduces malware via a USB drive.

The term risk represents the combination of the likelihood that a threat will exploit a vulnerability and the impact that the resulting incident would have on the organisation. In OT, risk is often expressed in terms of safety, production loss, financial cost, regulatory fines and reputational damage. A common way to visualise risk is the risk matrix, which plots likelihood on one axis and impact on the other, producing categories such as low, medium, high and critical. For example, the risk of a PLC firmware exploit may be rated as “high likelihood” because the vulnerability is publicly known, and “critical impact” because it could cause a plant

shutdown and safety breach.

A risk assessment is the structured activity of identifying threats, vulnerabilities and assets, and then estimating the likelihood and impact of each potential event. OT risk assessments typically follow the stages defined in IEC 62443-3-2, which include asset identification, threat modelling, vulnerability analysis and risk determination. The output is often a risk register, a living document that lists each identified risk, its current rating, the owner responsible for mitigation, and the planned treatment actions. The risk register enables traceability and supports decision-making by senior management.

The concept of likelihood refers to the probability that a specific threat will successfully exploit a given vulnerability. In OT, likelihood is influenced by factors such as network segmentation, the presence of intrusion detection systems, and the skill level of potential attackers. A quantitative approach might assign a numeric probability (e.G., 0.02 For a 2 % chance per year), while a qualitative approach uses descriptors like "rare", "possible" or "frequent". Selecting the appropriate method depends on data availability; many OT environments lack sufficient incident history for precise statistical modelling.

Impact describes the consequence of a risk event. In OT, impact is multidimensional: It can affect safety (injury or loss of life), environment (spills, emissions), production (downtime, lost output), financial performance (lost revenue, remediation cost), and compliance (breaches of the NIS Directive or industry-specific regulations). An impact analysis often involves estimating the monetary value of downtime (e.G., £500 000 Per hour) and the potential cost of a safety incident (e.G., £5 Million in compensation). The impact rating is a key driver for prioritising remediation efforts.

A risk matrix is a visual tool that helps stakeholders quickly understand which risks demand immediate attention. The matrix is typically a grid with likelihood categories on the vertical axis and impact categories on the horizontal axis. Risks that fall in the top-right quadrant (high likelihood, high impact) are treated as "critical". For example, a compromised SCADA server that could allow manipulation of valve positions would appear in the critical quadrant, prompting an urgent mitigation plan.

Residual risk is the amount of risk that remains after controls have been applied. In OT, eliminating risk entirely is rarely feasible because safety-critical systems must stay operational and some vulnerabilities cannot be fully remediated (e.G., Legacy hardware without vendor support). The goal is to reduce residual risk to a level that aligns with the organisation's risk appetite. For instance, after applying network segmentation, patching what is possible, and adding an intrusion prevention system, the residual risk might be downgraded from "critical" to "medium".

Risk appetite is the amount and type of risk that an organisation is willing to accept in pursuit of its objectives. It is set by senior leadership and expressed in policies that guide risk treatment decisions. An OT-focused enterprise may have a low appetite for safety-related risk but a higher appetite for financial risk if the cost of a mitigation outweighs the potential loss. The risk appetite statement should reference specific thresholds, such as "no more than one safety incident per 10 years".

Risk tolerance defines the acceptable deviation from the risk appetite. While risk appetite is a strategic stance, tolerance provides operational limits. For example, a company may tolerate a 5 % increase in downtime risk during a planned upgrade, but not a 20 % increase that would breach the risk appetite. Tolerance levels are useful for monitoring whether risk treatment actions are performing as expected.

Mitigation refers to the implementation of safeguards that reduce either the likelihood or the impact of a risk. In OT, mitigation strategies may include technical controls (firewalls, application whitelisting), procedural changes (regular backup of control logic), or organisational measures (security awareness training). A common mitigation is the deployment of a defence-in-depth architecture, where multiple layers of security (e.G., Perimeter firewalls, internal segmentation, host-based intrusion prevention) protect the same asset, making it harder for an attacker to succeed.

Control and safeguard are often used interchangeably; both denote a measure that reduces risk. Controls can be administrative (policies, training), technical (encryption, access control lists) or physical (locked cabinets, CCTV). In the OT domain, a technical control might be the use of a secure protocol such as IEC 60870-5-104 with TLS, while a physical safeguard could be a door-access system that prevents unauthorised entry to the control room.

Defence-in-depth is a design principle that advocates multiple, overlapping layers of security. In practice, this might involve separating the corporate IT network from the OT network with a demilitarised zone (DMZ), applying strict firewall rules, and then adding host-based intrusion detection on each PLC. The idea is that if one layer fails, the next layer still provides protection. The defence-in-depth concept aligns with IEC 62443, which recommends defence at the asset, zone, and system levels.

IEC 62443 is an internationally recognised series of standards for OT security. It provides a comprehensive framework covering policies, system design, risk assessment, and technical controls. The standard is divided into four parts: General concepts, policies and procedures, system design, and component requirements. For risk-management learners, IEC 62443-3-2 is particularly important because it outlines the process for risk assessment and the definition of security levels (SL-1 to SL-4). An example of applying IEC 62443 is using the standard to determine that a PLC controlling a high-pressure vessel must meet SL-3, which dictates specific authentication and encryption requirements.

NIST SP 800-82 is the US National Institute of Standards and Technology guide for OT security. Although it originates from the United States, the guidance is widely adopted in the United Kingdom as a complementary reference to IEC 62443. NIST SP 800-82 recommends a risk-based approach, the establishment of a security architecture, and the continuous monitoring of OT assets. It also provides a detailed taxonomy of threats, which can be mapped to the UK's NIS Directive requirements.

OT network segmentation is the practice of dividing the OT environment into logical zones to limit the spread of threats. Segmentation may be achieved using VLANs, firewalls, or physical air-gaps. A typical example is separating the engineering workstation network from the real-time control network, allowing

engineers to access historical data without being able to directly command field devices. Proper segmentation reduces the attack surface and simplifies the application of security policies.

PLC stands for programmable logic controller, a ruggedised computer used to control industrial processes. PLCs are often programmed with ladder logic, but they may also run more complex scripts. Because PLCs directly manipulate machinery, a compromise can have immediate safety consequences. Understanding the firmware version, communication ports and the authentication mechanisms of a PLC is essential for accurate risk evaluation.

SCADA (supervisory control and data acquisition) systems provide a high-level view of plant operations, aggregating data from many PLCs and allowing operators to issue commands. SCADA servers typically host historian databases, which store process data for analysis and reporting. Risks to SCADA include data integrity attacks (altering process values) and denial-of-service attacks that prevent operators from seeing real-time status.

DCS (distributed control system) is similar to SCADA but is often used in continuous processes such as chemical plants or refineries. DCS architectures are more tightly integrated, with control loops distributed across multiple controllers. Risk terminology for DCS includes “control loop integrity”, which refers to the assurance that the feedback loop is not tampered with.

HMI (human-machine interface) is the graphical interface through which operators interact with the control system. HMIs display process variables, alarms and trends, and may allow direct manipulation of set-points. A compromised HMI can be used to hide alarm floods or to issue deceptive commands, leading to unsafe states. Security controls for HMIs often include authentication, role-based access control and session time-outs.

Cyber-physical system (CPS) describes the integration of computation, networking and physical processes. In OT, every system is essentially a CPS, but the term is used to emphasise the tight coupling between the digital and physical worlds. Risk management for CPS must consider both cyber-threats and physical consequences, which can amplify each other. For example, a cyber attack that disables a pressure relief valve could cause a catastrophic physical explosion.

Safety Instrumented System (SIS) is a dedicated system that monitors hazardous conditions and initiates protective actions independent of the primary control system. SIS components are typically designed to meet IEC 61511 standards for functional safety. A key risk term is “safety integrity level” (SIL), which quantifies the reliability required for the SIS to perform its safety function. When conducting a risk assessment, analysts must verify that the SIS maintains its required SIL even in the presence of cyber threats.

Security lifecycle is the continuous process of planning, implementing, operating, monitoring and improving security controls. IEC 62443 defines a lifecycle that aligns with the engineering phases of OT: Design, implementation, operation, maintenance and decommissioning. Each phase includes risk-related activities such as threat modelling during design, vulnerability scanning during operation, and secure disposal

procedures at decommissioning.

Incident response in OT is the set of actions taken to detect, contain, eradicate and recover from a security incident. It differs from IT incident response because OT incidents often require coordination with safety teams, immediate shutdown procedures, and strict regulatory reporting. An incident response plan typically contains playbooks for scenarios like “unauthorised PLC programming change” or “malware infection of the historian”. Practising these playbooks through tabletop exercises helps ensure rapid, coordinated action.

Breach denotes a successful intrusion that results in unauthorised access to OT assets or data. Breaches can be classified by their scope (single device vs. Entire network) and by their effect (data exfiltration, process manipulation, safety impact). The 2022 ransomware incident at a UK water treatment facility is an illustrative breach: Attackers gained remote access, encrypted the SCADA server, and demanded payment, leading to a temporary shutdown of water supply.

Cyber threat intelligence (CTI) is the collection and analysis of information about current and emerging threats. In OT, CTI may include indicators of compromise (IOCs) specific to industrial protocols, reports on new vulnerabilities in PLC firmware, and geopolitical analysis of state-sponsored threat actors targeting critical infrastructure. Integrating CTI into the risk assessment process allows organisations to update likelihood scores as threat landscapes evolve.

Attack vector describes the path or method used by a threat to reach a vulnerability. Common OT attack vectors include phishing emails that deliver malicious attachments, compromised supply-chain components, insecure remote access tools, and malicious insiders with privileged credentials. Understanding the attack vector is essential for selecting appropriate mitigations; for example, if the primary vector is remote desktop protocol (RDP) abuse, then enforcing multi-factor authentication and limiting RDP exposure can drastically reduce risk.

Zero-day refers to a vulnerability that is unknown to the vendor and therefore has no patch available. Zero-day exploits are especially dangerous in OT because many devices cannot be patched quickly due to operational constraints. A well-known example is the 2017 vulnerability in certain Modbus-enabled PLCs that allowed unauthorised command injection before a patch was released. Managing zero-day risk involves compensating controls such as network isolation and strict access monitoring.

Phishing is a social-engineering technique used to trick users into revealing credentials or executing malicious code. In OT, phishing attacks often target engineering staff who have privileged access to control systems. A successful phishing email may deliver a credential-stealing payload that later logs into the OT network via a VPN. Mitigation includes regular security awareness training, email filtering, and the use of privileged-access management tools.

Insider threat encompasses both malicious and accidental actions by personnel with legitimate access. In OT, insiders may inadvertently introduce malware via removable media, or intentionally sabotage a process for personal gain. Risk assessments must consider insider threat because the insider already possesses the

knowledge and access required to bypass many technical controls. Countermeasures include segregation of duties, activity logging, and behavioural monitoring.

Supply-chain risk addresses the vulnerabilities that arise from third-party components, services or software. OT supply chains often involve equipment manufacturers, system integrators and software vendors. A compromised firmware update from a component supplier can introduce a backdoor into a PLC. Managing supply-chain risk requires due-diligence, verification of vendor security practices, and the use of cryptographic signing of firmware.

Third-party risk is similar to supply-chain risk but focuses on services such as cloud-based monitoring, managed security services, and remote support contracts. Third-party risk assessments evaluate the security posture of service providers, their incident-response capabilities, and contractual obligations for data protection. For example, a cloud-based analytics platform that stores process data must be assessed for confidentiality and availability guarantees.

Compliance refers to adherence to legal, regulatory and industry standards. In the United Kingdom, OT risk management must align with the NIS Directive (Network and Information Systems), the UK Data Protection Act (which incorporates GDPR), and sector-specific guidance such as the UK Water Industry's "Cyber Security Guidance". Non-compliance can result in fines, enforcement actions and loss of licence.

GDPR (General Data Protection Regulation) applies to personal data, which may be present in OT environments through employee records, contractor details or customer information stored in SCADA databases. While GDPR primarily concerns privacy, its breach notification requirements intersect with OT incident reporting. For instance, a ransomware event that encrypts a historian containing personal data would trigger a GDPR breach notice within 72 hours.

NIS Directive (now the UK's NIS Regulations) mandates that operators of essential services implement appropriate security measures and report incidents. The directive defines "appropriate and proportionate" measures, which are interpreted through risk assessments. Failure to meet NIS obligations can lead to civil penalties and increased scrutiny from the regulator.

ISO 27001 is an international standard for information security management systems (ISMS). Although it is IT-centric, many OT organisations adopt ISO 27001 to demonstrate a systematic approach to risk management, policy development and continual improvement. The standard's Annex A provides a set of controls that can be mapped to OT-specific safeguards.

ISO 31000 provides principles and guidelines for risk management across all sectors. It emphasises the creation of a risk-aware culture, the integration of risk processes into decision-making, and the need for risk communication. OT risk practitioners often reference ISO 31000 when defining the organisational framework for risk governance.

Risk analysis is the analytical phase of risk assessment where likelihood and impact are quantified or

qualified. Qualitative risk analysis uses descriptive scales (e.G., "High", "medium"), while quantitative risk analysis employs numerical values, probability distributions and statistical models. Quantitative analysis can produce a monetary risk value (e.G., £2 Million per year), which is useful for cost-benefit calculations of mitigation measures.

Qualitative vs quantitative approaches each have advantages. Qualitative methods are faster, require less data, and are useful when historical incident data is scarce. Quantitative methods provide more precise estimates but demand extensive data collection, such as failure rates of components, cost of downtime, and frequency of threat events. Many OT risk assessments use a hybrid approach: They start with qualitative scoring to prioritise risks, then apply quantitative analysis to the highest-priority items.

Risk scoring assigns a numeric value to each risk based on the product of likelihood and impact scores. A common formula is $\text{Risk Score} = \text{Likelihood} \times \text{Impact}$. Scores enable ranking of risks and support the selection of treatment options. For example, a risk with a likelihood of 4 (on a 1-5 scale) and an impact of 5 would receive a score of 20, placing it in the top tier for remediation.

Risk heat map is a visual representation similar to the risk matrix but often colour-coded to highlight "hot" (high-risk) areas. Heat maps are useful for executive briefings because they provide a quick visual of the overall risk posture. A heat map might show clusters of risk in the "remote access" and "legacy device" categories, signalling where strategic investment is needed.

Risk owner is the individual or function accountable for managing a given risk. In OT, risk owners are typically system engineers, plant managers or the head of operations. The risk owner is responsible for ensuring that mitigation actions are implemented, that the risk rating is reviewed regularly, and that any changes in the environment are reflected in the risk register. Clear ownership prevents "risk drift", where responsibilities become unclear over time.

Risk treatment encompasses the four primary strategies: Avoidance, reduction, sharing and acceptance. In OT, avoidance might involve decommissioning an obsolete PLC that cannot be patched. Reduction includes applying technical controls such as firewalls. Sharing could be achieved through cyber-insurance, which transfers part of the financial impact to an insurer. Acceptance occurs when the residual risk is deemed tolerable relative to the organisation's appetite.

Risk acceptance is the conscious decision to retain a risk because the cost of mitigation exceeds the benefit, or because the risk falls within the defined tolerance. Acceptance must be documented, approved by senior management, and communicated to all stakeholders. An example is accepting the risk of a non-critical HMI that cannot be upgraded due to compatibility issues, after confirming that it does not influence safety functions.

Risk transfer involves shifting the financial consequences of a risk to another party, commonly through insurance contracts or outsourcing agreements. In OT, cyber-insurance policies may cover costs associated with incident response, legal fees and business interruption. However, insurers often require demonstrable

security controls, so risk transfer can also act as a catalyst for improving the underlying security posture.

Risk avoidance is the elimination of a risk by removing the source of the threat or the vulnerable asset. For instance, an organisation may avoid the risk of a vulnerable field-bus protocol by replacing it with a secure Ethernet-based solution. While avoidance can be effective, it may not always be practical due to cost, downtime or regulatory constraints.

Risk sharing is a collaborative approach where multiple parties jointly manage a risk. In OT, joint-venture projects may share the responsibility for securing a shared water-treatment plant, using a common risk-management framework. Formal agreements define each party's obligations, liability limits and communication channels.

Business continuity (BC) is the capability of an organisation to continue essential functions during and after a disruptive event. In OT, BC plans encompass scenarios such as loss of network connectivity, power outages, or cyber attacks. The BC plan outlines alternative operating modes, manual overrides and recovery procedures to maintain production and safety.

Disaster recovery (DR) focuses on restoring IT and OT systems after a catastrophic event. DR plans specify recovery objectives, data restoration processes, and the sequencing of system bring-up. In OT, DR must consider the order of control system restoration to avoid unsafe states; for example, the safety-instrumented system must be verified before the production control system is re-energised.

Recovery Time Objective (RTO) is the maximum acceptable duration between a disruption and the restoration of a specific service. For a critical PLC controlling a reactor, an RTO of 30 minutes may be required to prevent safety degradation. RTO values guide the design of redundancy, failover mechanisms, and spare parts inventory.

Recovery Point Objective (RPO) defines the maximum tolerable period of data loss measured in time. In a SCADA historian, an RPO of 5 minutes means that data older than five minutes may be lost without unacceptable impact. RPO influences backup frequency, replication strategies and the selection of storage technologies.

Resilience is the ability of a system to absorb disturbances and continue operating. In OT, resilience is achieved through design principles such as redundancy, diversity, and graceful degradation. A resilient control system might automatically switch to a backup PLC if the primary controller fails, while maintaining safe operation.

Redundancy involves duplicating critical components so that a failure of one does not affect overall functionality. Redundancy can be active (both units operate simultaneously) or standby (the backup is idle until needed). In a power-generation plant, redundant communication paths between the DCS and field devices ensure that a single cable cut does not halt control.

Failover is the automatic transfer of control from a failed component to its redundant counterpart. Successful failover requires health monitoring, rapid detection of failure, and a predefined switchover procedure. An example is a redundant PLC pair that synchronises state information; when the primary PLC loses heartbeat, the secondary assumes control within seconds.

Hardening refers to the process of reducing the attack surface of an asset by disabling unnecessary services, applying patches, and enforcing strong authentication. Hardening guidelines for OT devices often come from manufacturers or from standards such as IEC 62443-4-2. A hardened PLC might have only the required Modbus port open, with all other services blocked, and use certificate-based authentication.

Patch management is the systematic process of acquiring, testing, and deploying software updates. In OT, patch management is challenging because many devices cannot be taken offline without impacting production, and vendor support periods may be limited. A typical approach is to schedule patch windows during planned outages, conduct thorough regression testing, and maintain a rollback plan.

Change management governs the process of modifying OT systems, ensuring that changes are reviewed, approved, documented and tested before implementation. Effective change management reduces the risk of accidental mis-configurations that could introduce new vulnerabilities. A change request for updating a PLC program, for example, must include a risk assessment, a back-out plan and sign-off from the safety engineer.

Configuration management maintains the integrity of system settings over time. It involves maintaining a baseline configuration, tracking deviations, and enforcing authorised configurations. In OT, configuration drift can occur when field engineers manually edit device settings without updating the central repository, leading to inconsistencies that increase risk.

Asset inventory is the comprehensive catalogue of all hardware and software components within the OT environment. An accurate inventory is essential for vulnerability scanning, compliance reporting and risk prioritisation. Asset inventories often integrate with CMDBs (configuration management databases) and can be populated automatically using network discovery tools that recognise industrial protocols.

Baseline defines the normal or expected state of an asset, against which deviations are measured. Baselines may include operating system versions, open ports, and permitted communication flows. Establishing a secure baseline enables rapid detection of unauthorised changes, supporting both compliance and incident response.

Anomaly detection involves monitoring system behaviour to identify deviations from the baseline that may indicate a security event. In OT, anomaly detection can be protocol-aware, analysing Modbus traffic patterns for unexpected function codes. Machine-learning models are increasingly used to detect subtle changes that traditional signature-based tools might miss.

Intrusion Detection System (IDS) monitors network traffic for signs of malicious activity. In OT, IDS solutions

are often protocol-specific, able to understand the semantics of industrial traffic. An IDS may generate alerts when it observes a command that attempts to write to a protected register on a PLC, signalling a potential intrusion.

Intrusion Prevention System (IPS) extends IDS functionality by actively blocking malicious traffic. Deploying an IPS in an OT network must be done carefully to avoid unintended disruption of legitimate control traffic. Policy tuning is crucial; for example, an IPS rule that blocks all Modbus write commands could prevent normal operations and cause false positives.

Security Operations Center (SOC) is a dedicated team that monitors security events, analyses alerts and coordinates incident response. In many organisations, the SOC is IT-centric, but for OT it must include specialists who understand industrial protocols and safety implications. A SOC may ingest data from OT-specific SIEMs, IDS/IPS sensors, and physical security systems to provide a holistic view.

Security Information and Event Management (SIEM) aggregates logs from diverse sources, correlates events, and provides dashboards for analysts. OT-focused SIEMs ingest data from PLCs, SCADA servers, firewalls, and physical access controls. Configuring meaningful correlation rules is essential; for instance, a rule that correlates a successful VPN login with a subsequent PLC configuration change can highlight a suspicious activity chain.

Threat hunting is the proactive search for indicators of compromise that have evaded automated detection. In OT, threat hunting may involve reviewing command histories, analysing firmware version anomalies, or inspecting unusual network flows between engineering workstations and field devices. Successful threat hunting can uncover hidden backdoors or dormant malware before they cause damage.

Security policy is a formal document that defines the organisation's security objectives, roles, responsibilities and acceptable behaviours. OT security policies must address both cyber and safety aspects, specifying requirements for network segmentation, access control, patch management, and incident reporting. Policies are often referenced during audits and serve as the basis for employee training.

Governance refers to the framework of rules, practices and processes that ensure the achievement of an organisation's security objectives. Effective governance links risk management activities with strategic business goals and regulatory requirements. In OT, governance structures often include a dedicated cyber-security steering committee that reviews risk assessments, monitors compliance, and allocates resources.

Audit is an independent examination of processes, controls and compliance with policies. OT audits may be internal or external, and can focus on technical controls (e.G., Firewall rule reviews), procedural compliance (e.G., Change-management records) or regulatory adherence (e.G., NIS reporting). Audits provide assurance that risk-management processes are being followed and highlight gaps for remediation.

Penetration testing (pen-test) simulates real-world attacks to evaluate the effectiveness of security controls.

In OT, pen-testing must be carefully scoped to avoid disrupting critical processes. A common approach is to conduct “red-team” exercises on a test environment that mirrors the production network, or to use passive techniques such as vulnerability scanning without active exploitation.

Red team is a group that adopts the attacker’s perspective, attempting to breach security controls using realistic tactics. Red-team engagements in OT can reveal hidden pathways, such as the use of a legacy remote-access tool that provides unauthorised entry to the control network. Findings from red-team activities feed directly into risk-treatment planning.

Blue team defends the environment, monitors alerts, and responds to incidents. In an OT context, the blue team includes control-system engineers, safety officers and SOC analysts. Collaboration between red and blue teams during joint exercises improves detection capabilities, refines response procedures and strengthens overall security posture.

Cyber insurance provides financial protection against losses resulting from cyber incidents. Policies may cover costs related to incident response, legal liability, regulatory fines, and business interruption. When underwriting OT cyber-insurance, insurers often require evidence of robust risk-management practices, such as documented risk assessments, security controls aligned with IEC 62443, and regular testing.

Incident handling is the systematic process of managing a security incident from detection through resolution. It includes phases such as preparation, identification, containment, eradication, recovery and lessons-learned. In OT, incident handling must coordinate with safety teams to assess any immediate physical hazards and may involve shutting down equipment to prevent escalation.

Root cause analysis (RCA) investigates the underlying reasons for an incident, beyond the immediate trigger. RCA techniques such as the “5 Whys” or fishbone diagrams help uncover systemic weaknesses, like inadequate access controls or insufficient monitoring. Conducting RCA after an OT incident is essential for updating the risk register and preventing recurrence.

Threat modelling is the practice of systematically enumerating potential threats, their motivations, capabilities and attack paths. In OT, threat modelling often uses methodologies such as STRIDE (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege) adapted for industrial protocols. The output is a set of threat scenarios that feed into likelihood assessments.

Attack surface denotes all points where an adversary could attempt to gain unauthorised access. Reducing the attack surface is a key mitigation strategy; for OT this may involve disabling unused communication ports on PLCs, removing legacy protocols, and limiting remote access to a single hardened jump host.

Security architecture is the high-level design that defines how security controls are organised and integrated. A well-designed OT security architecture incorporates defence-in-depth layers, network segmentation, identity and access management, and monitoring components. Architectural diagrams often illustrate zones (e.G., Enterprise, DMZ, control, safety) and the flow of data between them.

Identity and Access Management (IAM) governs who can access which resources, under what conditions. In OT, IAM must support role-based access control (RBAC) that distinguishes between operators, engineers, maintenance staff and auditors. Multi-factor authentication (MFA) is increasingly required for remote access to control-system consoles.

Role-based access control (RBAC) assigns permissions based on job function rather than individual identity. RBAC reduces the risk of privilege creep, where users accumulate unnecessary rights over time. For example, an operator may have read-only access to process data, while an engineer has write permissions to modify PLC logic, and a safety officer has authority to trigger emergency shutdowns.

Multi-factor authentication (MFA) requires two or more verification methods, such as a password and a hardware token. MFA significantly lowers the probability of credential compromise leading to unauthorised access. In OT, MFA is commonly applied to VPN gateways, jump servers and web-based HMI portals.

Least privilege is the principle that users, processes and systems should be granted only the minimum permissions necessary to perform their functions. Implementing least privilege in OT reduces the impact of a compromised account. For instance, a maintenance technician's account might be limited to rebooting devices, without the ability to alter control logic.

Segregation of duties (SoD) separates critical functions among multiple individuals to prevent fraud or error. In OT, SoD might mean that the person who approves a change is not the same person who implements it. SoD controls are documented in the risk register and verified during audits.

Network monitoring involves continuous observation of traffic flows, device health and security events. Effective network monitoring in OT requires protocol-aware sensors that can interpret Modbus, OPC UA, EtherNet/IP and other industrial traffic. Monitoring enables early detection of anomalies such as unexpected command sequences or abnormal data rates.

Physical security protects the tangible assets of an OT environment. It includes measures such as perimeter fencing, access cards, biometrics, CCTV and security patrols. Physical security is integral to risk management because unauthorised physical access can lead to direct manipulation of control devices or the installation of rogue hardware.

Environmental monitoring tracks conditions such as temperature, humidity and vibration that could affect equipment reliability. While not a traditional security control, environmental monitoring contributes to risk management by identifying conditions that may cause premature failure, which in turn can create safety or production risks.

Supply chain integrity ensures that hardware and software components are authentic and have not been tampered with before deployment. Techniques such as cryptographic signing of firmware, secure boot, and hardware-based root of trust help verify integrity. Verifying supply-chain integrity is especially important for devices sourced from overseas manufacturers.

Secure boot validates the integrity of firmware during system start-up, preventing unauthorised code from executing. Many modern PLCs support secure boot, which checks digital signatures against a trusted key stored in hardware. Deploying secure boot reduces the risk of firmware-level attacks that could persist across reboots.

Cryptographic signing uses digital certificates to verify that software or firmware originates from a trusted source. In OT, signed firmware updates provide assurance that the code has not been altered in transit. Establishing a robust PKI (public key infrastructure) is essential for managing certificates and revocation lists.

Public key infrastructure (PKI) is the framework for creating, distributing, managing and revoking digital certificates. PKI enables secure communications, authentication and code signing. OT deployments often use PKI for TLS encryption of OPC UA traffic, VPN authentication, and firmware signing.

TLS (Transport Layer Security) encrypts data in transit, protecting confidentiality and integrity. When applied to industrial protocols such as OPC UA, TLS prevents eavesdropping and man-in-the-middle attacks.