

Rehan Joseph ACCA5034

by Rehan Joseph

Submission date: 05-Jan-2026 11:12PM (UTC+0000)

Submission ID: 272210690

File name: 57290_Rehan_Joseph_Rehan_Joseph_ACCA5034_712345_614830875.docx (632.3K)

Word count: 6720

Character count: 44874

REHAN JOSEPH


(2404834)

Data Security and Compliance

ACCA5034

Assignment

6 January 2026



Contents

Figures	3
Table of Acronyms.....	3
Task 1	5
Introduction	5
Why AI Governance Is Important	5
AI as a High-Impact but High-Risk Technology	5
Legal and Ethical Responsibilities.....	6
Governance as Protection and Enablement.....	6
Anticipating Risks in AI Systems	6
Risk Assessment Before Development	7
Data Quality, Bias, and Fairness.....	7
Ethical and Social Impact.....	7
Anticipating Security Threats	7
Building a Preventative Governance Mindset	8
Protecting AI Systems Throughout Their Lifecycle	8
Robust Data Governance	8
Secure Model Development Practices	8
Ethical Frameworks and Embedded Principles	8
Legal and Regulatory Compliance.....	9
Detecting Problems in AI Systems	9
Responding to AI Incidents in Real Time.....	9
Embedding Response Within a Comprehensive AI Governance Framework	10
Case Examples	10
Conclusion.....	11
TASK 2	12
1. Introduction	12
2. Organisational Structure: Embedding Zero Trust into People and Processes	13
3. Policies and Procedures: Translating Zero Trust into Everyday Practice	14
AI Usage Policy	15
Secure SaaS Configuration Policy	15
BYOD Policy	15

4. Regulatory Compliance: Addressing DORA and NIS2 Requirements 16

5. Governance and Strategy: Aligning Security with Business Goals Using NIST CSF 2.0
..... 17

6. Risk Management: Focus on AI and Cloud Threats 19

7. Contingency Planning: Ransomware Impacting AI Assets 20

8. Protection Mechanisms: Enforcing Zero Trust Through Modern Controls..... 20

9. Conclusion 21

Figures

Figure 1: AI Governance as a balancing mechanism 6

Figure 2: Summary of Case Studies 11

Figure 3: Organisational Structure 14

Figure 4 Policy Table 16

Figure 5 : NIST Cybersecurity Framework (CSF) 17

Figure 6 19

Table of Acronyms

Acronym	Full Term
AI	Artificial Intelligence
ML	Machine Learning
BYOD	Bring Your Own Device
CI/CD	Continuous Integration / Continuous Deployment
CSPM	Cloud Security Posture Management
DORA	Digital Operational Resilience Act
EDR	Endpoint Detection and Response
GDPR	General Data Protection Regulation
IAM	Identity and Access Management
MFA	Multi-Factor Authentication
NIS2	Network and Information Systems Directive 2
NIST	National Institute of Standards and Technology
SaaS	Software as a Service

SIEM	Security Information and Event Management
SOAR	Security Orchestration, Automation, and Response
ZTNA	Zero Trust Network Access

Task 1

Governance of AI and Machine Learning Systems: Managing risk, ethics, and compliance for AI data and models.

Introduction

Artificial Intelligence (AI) and Machine Learning (ML) are no longer futuristic concepts; they are now important to modern organizations[1]. Banks use AI to detect fraud, hospitals use ML to accelerate diagnoses, schools analyse data to predict student performance, and transport systems increasingly rely on semi-autonomous technologies. AI shapes how decisions are made, how services are delivered, and how organizations compete.

The rapid adoption of AI has delivered significant benefits, including faster decision-making, increased efficiency, personalized services, and insights that would be impossible for humans to uncover alone[2]. However, these advantages are accompanied by serious risks. Unlike traditional software, AI systems do not operate solely on fixed rules; they learn from data, adapt over time, and may behave in unpredictable ways[3].

As a result, AI systems can reinforce existing biases, compromise privacy, or be manipulated by malicious actors. When failures occur, their impact can scale rapidly, affecting thousands or even millions of individuals. Such incidents can lead to financial loss, legal consequences, physical harm, and long-term reputational damage[3].

To address these challenges, organizations require robust AI governance. AI governance includes the policies, ethical principles, technical safeguards, and oversight mechanisms that ensure AI systems are trustworthy, secure, and compliant with legal and regulatory requirements. Effective governance enables organizations to use the

benefits of AI while maintaining accountability, transparency, and resilience[4].

This paper examines how organizations can implement AI governance frameworks that anticipate risks, protect AI systems throughout their lifecycle, detect issues in real time, and respond effectively to incidents. Through analysis and real-world examples, it demonstrates how strong governance is essential for managing risk, ensuring ethical use, and maintaining trust in an increasingly AI-driven world.

Why AI Governance Is Important

AI as a High-Impact but High-Risk Technology

AI systems are being used more often in decision-making processes that directly affect individuals and society[5]. Their ability to process vast datasets, identify complex patterns, and operate at scale allows organizations to reduce costs, increase efficiency, and gain competitive advantage. However, as AI moves from decision support to decision automation, the risks associated with errors, bias, and misuse increase significantly. Without structured governance, these systems are often deployed with insufficient validation, oversight, or ethical review, increasing the likelihood of harmful outcomes[6].

Unlike traditional software systems, AI models can produce outcomes that are difficult to predict or explain, particularly in the case of complex machine learning and deep learning models. These systems may influence critical outcomes such as credit approval, employment screening, medical diagnosis, or criminal risk assessment. When such systems produce biased, inaccurate, or unclear decisions, the consequences can include discrimination, financial loss, physical harm, and violations of fundamental rights[6]. Because AI systems

operate at scale, a single flawed model can cause widespread harm in a very short time.

The lack of clarity of many AI models further complicates accountability. When decision logic cannot be clearly explained, it becomes difficult to challenge outcomes, assign responsibility, or demonstrate compliance[6]. This creates risks not only for affected individuals but also for organizations that deploy these systems without adequate oversight.

Legal and Ethical Responsibilities

In response to these risks, regulators are introducing detailed AI-specific laws. The European Union's AI Act adopts a risk-based framework that imposes strict requirements on high-risk AI systems, including obligations related to data quality, transparency, human oversight, and ongoing risk management[5]. Similar regulatory approaches are emerging in the UK, the United States, and other countries, signalling a global shift toward stricter AI accountability. However, regulatory requirements vary widely across countries, making compliance more difficult for multinational organizations and increasing the likelihood of uneven governance practices.

Failure to comply with these regulations can result in substantial financial penalties, operational restrictions, and reputational damage[5]. However, legal compliance alone is insufficient. Organizations are increasingly expected to demonstrate ethical responsibility by ensuring fairness, transparency, privacy protection, and harm prevention. Public trust in AI systems depends not only on legality but on perceived legitimacy and accountability[5].

AI governance provides the mechanisms through which these legal and ethical requirements can be put into practice. By embedding responsibility into model development, deployment, and monitoring processes, governance frameworks ensure that ethical principles are translated into

enforceable practices rather than aspirational statements[5].

Governance as Protection and Enablement

Effective AI governance acts as both a protective and enabling force. It protects individuals by reducing the likelihood of unfair or harmful outcomes, and it protects organizations by mitigating legal, financial, and operational risks. Governance frameworks establish clear roles, decision rights, and escalation paths, enabling faster and more effective responses to incidents when failures occur[4].

Crucially, strong governance does not inhibit innovation. Instead, it provides the clarity and confidence required to deploy AI systems responsibly at scale. Organizations with mature governance structures are better positioned to innovate safely, adapt to regulatory change, and maintain trust with customers, employees, and regulators[4].




Figure 1: AI Governance as a balancing mechanism

Figure 1 illustrates how AI governance functions as a balancing mechanism between the benefits of AI adoption and the associated ethical, legal, and operational risks. It shows governance as an enabling control that allows organisations to realise efficiency, innovation, and improved decision-making while mitigating risks such as bias, privacy loss, and security threats.

Anticipating Risks in AI Systems

Anticipating risk is a foundational element of AI governance. It involves identifying, assessing,

and mitigating potential harms *before* an AI system is deployed into real-world environments. This stage is critical because many AI-related harms such as discriminatory outcomes, loss of privacy, or erosion of trust are difficult or impossible to reverse once systems are operational and embedded in organisational processes.

Risk Assessment Before Development

Effective AI governance requires structured risk assessments prior to system development and deployment. These assessments evaluate technical, ethical, legal, and societal risks associated with a proposed AI system[6]. Key considerations include the nature of the decisions being automated, the populations affected, and the potential consequences of errors, bias, or misuse.

For example, an AI system used for loan approval operates in a high-impact context, as incorrect or biased decisions can directly affect individuals' financial stability and access to opportunities. Governance-driven risk assessments help organisations determine whether AI is appropriate for a given task and identify necessary safeguards, such as human oversight, explainability requirements, and accessible appeal mechanisms[6]. Without this early-stage scrutiny, risks may be embedded into system design and spread at scale.

Data Quality, Bias, and Fairness

Data governance is central to anticipating AI risk. Machine learning models learn patterns from historical data, which may reflect existing social inequalities, institutional biases, or outdated assumptions. If these issues are not addressed during data selection and preparation, AI systems can repeat and increase unfair outcomes[7].

Governance frameworks require organisations to assess datasets for accuracy, representativeness, and relevance to the intended use case. This includes identifying

indirect indicators that may indirectly encode sensitive attributes such as gender, ethnicity, or socio-economic status. Techniques such as bias audits, fairness testing, and preparing data are essential governance controls that reduce discriminatory risk and improve system reliability[7].

Ethical and Social Impact

Anticipating risk also extends beyond technical considerations to the broader ethical and social implications of AI deployment. AI systems can influence behaviour, restrict autonomy, and reinforce existing power imbalances[6]. Ethical impact assessments examine how systems affect fundamental values such as fairness, dignity, and equality, as well as their potential impact on public trust.

The use of predictive policing technologies illustrates these concerns. Such systems have been widely criticised for unfairly targeting already over-policed communities, thereby reinforcing cycles of surveillance and discrimination[8]. Anticipating these impacts allows organisations to redesign systems, limit their scope, introduce additional safeguards, or avoid deployment altogether in sensitive contexts.

Anticipating Security Threats

AI systems introduce distinct security risks that must be addressed as part of governance planning[6]. Threats such as adversarial attacks, data poisoning, and model extraction can compromise system integrity, confidentiality, and reliability. These risks are often overlooked until after deployment, when exploitation can cause widespread harm.

By anticipating security threats early in the lifecycle, organisations can implement preventative controls, including adversarial testing, secure data pipelines, model access restrictions, and continuous vulnerability assessment. Integrating security considerations

into AI governance reduces the likelihood of system compromise and strengthens overall resilience[6].

Building a Preventative Governance Mindset

A preventative approach to AI governance prioritises foresight over reaction. Rather than responding to incidents after harm has occurred, organisations embed risk anticipation into decision-making, development processes, and organisational culture. This proactive mindset is essential for ensuring that AI systems are safe, fair, and trustworthy from the outset, and it forms the basis for sustainable and responsible AI adoption

Protecting AI Systems Throughout Their Lifecycle

Protecting AI systems requires the implementation of safeguards that extend across the entire system lifecycle, from data collection and model development to deployment, operation, and ongoing maintenance. Unlike traditional software, AI systems continue to learn, interact with new data, and operate in changing environments, meaning that risks can emerge or evolve long after initial deployment. Effective AI governance therefore treats protection as a continuous responsibility rather than a one-time technical exercise[9].

Robust Data Governance

Data governance is a critical component of AI system protection, as data underpins model behaviour, performance, and fairness.

Governance frameworks require organisations to manage data responsibly through measures such as strict access controls, encryption of data at rest and in transit, the masking of personal information, and data minimisation practices that limit unnecessary exposure[9].

These controls reduce the risk of data breaches, support compliance with privacy regulations such as the GDPR, and protect individuals from misuse of their personal information[9]. In high-risk domains such as healthcare or finance, failures in data governance can lead to serious harm and rapid loss of public trust, reinforcing the need for strong, enforceable controls throughout the data lifecycle.

Secure Model Development Practices

AI governance also extends to how models are developed, tested, and maintained. Secure development practices include the use of version control systems to track changes, regular testing and validation to identify errors or unintended behaviour, and peer review processes to improve model quality and robustness[9]. Comprehensive documentation of design decisions, assumptions, training data, and known limitations further supports transparency and accountability[9].

These practices not only reduce the likelihood of vulnerabilities and failures but also enable organisations to explain and justify system behaviour when decisions are challenged by regulators, auditors, or affected individuals.

Ethical Frameworks and Embedded Principles

Protecting AI systems involves embedding ethical principles such as fairness, transparency, accountability, privacy, and safety directly into organisational workflows. Ethical frameworks guide decision-making when trade-offs arise between performance, efficiency, and risk, ensuring that ethical considerations are addressed alongside technical and business objectives[10].

By integrating ethical principles into development standards, review processes, and approval gates, organisations move beyond abstract values and establish consistent, enforceable expectations across teams and projects[10]. This consistency is essential for

maintaining trust and avoiding fragmented or ad-hoc ethical decision-making.

Legal and Regulatory Compliance

Lifecycle protection is also necessary to demonstrate ongoing compliance with AI-related laws, standards, and regulatory frameworks. Instruments such as the EU AI Act, the NIST AI Risk Management Framework, and relevant ISO standards provide structured guidance for managing AI risks across development and deployment stages.

Governance mechanisms that track compliance over time help organisations respond to regulatory scrutiny and adapt to evolving legal requirements[10].

Effective protection across the lifecycle enhances system resilience, enabling AI systems to operate safely, reliably, and responsibly in real-world environments.

Detecting Problems in AI Systems

Even with robust preventative controls and lifecycle protections in place, AI systems remain vulnerable to degradation, failure, and unintended behaviour once deployed. Changes in data distributions, evolving user behaviour, and emerging security threats can all undermine model performance over time. Effective AI governance therefore requires continuous detection mechanisms that identify problems early and enable timely intervention before harm escalates.

Continuous monitoring is a core governance control for deployed AI systems. Organisations must track not only technical performance metrics such as accuracy and latency, but also indicators related to fairness, stability, and reliability[9]. Model drift can occur when real-world data no longer reflects training conditions, leading to biased or inaccurate outcomes. Early detection of such drift allows organisations to

initiate corrective actions, including retraining, recalibration, or temporary suspension of system use, thereby limiting the spread of harm.

Detection also plays a critical role in protecting AI systems from malicious activity. Monitoring tools can identify unusual input patterns, repeated probing attempts, or sudden behavioural changes that may signal adversarial attacks, data poisoning, or unauthorised access[9]. Rapid identification of security threats helps preserve system integrity and reinforces trust among users, regulators, and stakeholders.

Auditability and traceability further strengthen detection capabilities. Comprehensive logging of model decisions, data usage, and system access enables organisations to investigate incidents, demonstrate accountability, and support regulatory audits. In high-risk contexts, such transparency is essential for understanding how and why failures occur, as well as for proving compliance with legal and ethical requirements.

Finally, detection mechanisms must extend beyond technical monitoring to include ongoing ethical and legal review. As regulations, societal expectations, and organisational values evolve, AI systems that were once compliant may no longer meet required standards. Regular reviews ensure that systems remain aligned with current legal obligations and ethical norms, allowing governance frameworks to adapt proactively rather than reactively.

Responding to AI Incidents in Real Time

Detecting issues in AI systems is only effective if organisations can respond quickly and decisively. Given the speed and scale at which AI systems operate, delayed responses can allow errors, bias, or security breaches to spread rapidly. Effective AI governance therefore requires response mechanisms that combine automated controls with clearly defined human oversight.

Automated response capabilities play a critical role in containing harm. Features such as model rollback allow organisations to revert to previously approved and validated model versions when performance degrades or unintended behaviour is detected[9]. This limits the operational impact of failures and prevents flawed models from continuing to influence decisions. Similarly, real-time alerting systems enable rapid escalation to engineers, security teams, compliance officers, and management when predefined thresholds are breached, ensuring accountability and informed intervention.

MLOps frameworks further support effective response by automating retraining, deployment, validation, and monitoring processes. Automation improves consistency, reduces human error, and allows systems to adapt to changing data or environments while remaining aligned with governance controls. In cases of suspected compromise or severe malfunction, isolation and containment mechanisms can automatically restrict system access or suspend operation, limiting the spread of harm while investigations and remediation take place[9].

Embedding Response Within a Comprehensive AI Governance Framework

While automation is essential, it is not sufficient on its own. A comprehensive AI governance framework ensures that response mechanisms operate within clearly defined organisational structures, policies, and accountability arrangements[10]. Governance provides the context that determines when systems may act autonomously, when human intervention is required, and who is ultimately responsible for decisions and outcomes.

Clear policies and standards define acceptable AI use, risk tolerance, escalation procedures, and compliance requirements. These policies ensure consistent governance across large or complex organisations and prevent fragmented or ad-hoc responses to incidents. Defined

oversight roles such as AI ethics committees, governance boards, and data protection officers maintain accountability and ensure that high-impact decisions remain subject to human judgement[10].

Training and organisational awareness are also critical. Employees involved in the development, deployment, and oversight of AI systems must understand governance requirements, ethical responsibilities, and response procedures. A culture that values transparency, responsibility, and continuous improvement enable organisations to prioritise safety and trust over short-term performance gains, strengthening long-term resilience.

Case Examples

Real-world applications of AI decision-support systems show that ethical outcomes depend strongly on how these systems are designed, implemented, and governed. As summarised in the **figure 2**, well-governed AI deployments can deliver significant benefits, whereas poorly governed systems have led to serious ethical failures, including discrimination and breaches of privacy[6].

The first case study, **Ethical AI in Healthcare**, demonstrates how AI can be deployed responsibly to support clinical decision-making. Diagnostic tools used for detecting conditions such as cancer and cardiovascular disease have improved accuracy and efficiency while maintaining essential human oversight. By positioning AI as an assistive technology and ensuring that medical professionals retain final decision-making authority, these systems reduce the risks associated with full automation and uphold key ethical principles such as accountability and patient safety[6].

In contrast, figure 2 also highlights examples of unethical AI use within organizations. The **AI Bias in Hiring** case shows how recruitment algorithms have reproduced gender and racial biases when trained on historically biased hiring data, resulting in discriminatory outcomes rather than promoting fairness. Similarly, the

Discriminatory Learning Algorithms case illustrates how lending systems in financial services have disproportionately disadvantaged low-income and minority applicants by relying on biased historical decision data. These cases highlight the role of governance mechanisms such as bias audits, human oversight requirements, and accountability structures in determining whether AI systems produce ethical or harmful outcomes[6].

Case Study	Description
Ethical AI in Healthcare	AI systems are used to support medical professionals by improving diagnostic accuracy while retaining human oversight and adhering to ethical standards.
AI Bias in Hiring	Recruitment AI has reinforced gender and racial bias due to training on historically biased data, leading to discriminatory hiring outcomes.
Discriminatory Learning Algorithms	Lending AI systems have perpetuated socioeconomic and racial discrimination by relying on biased historical decision data.

Figure 2: Summary of Case Studies

Conclusion

AI governance is no longer optional; it is essential for modern organizations operating in data-driven and automated environments. Effective governance ensures AI systems remain safe, ethical, compliant, and reliable throughout their lifecycle. By anticipating risks, protecting systems, detecting issues early, and responding quickly, organizations can minimize harm while maximizing AI value.

Strong governance turns AI from a potential risk into a trusted strategic tool. It enables accountability, regulatory compliance, and public confidence. Rather than limiting innovation, governance provides the structure and assurance needed to use AI responsibly at scale.

Ultimately, the long-term success of AI depends not only on technical capability but also on strong oversight, ethical commitment, and organizational responsibility. Organizations that make governance a core part of their AI strategy will be best positioned to sustain trust, adapt to changing regulations, and succeed in an increasingly complex digital world.

TASK 2

Zero Trust–Aligned Security Programme for ResilientIntellect Ltd (EU Expansion)

1. Introduction

ResilientIntellect Ltd is a rapidly expanding threat intelligence and consultancy organisation specialising in AI-powered threat analysis software and high-impact cyber advisory services. Its fully hybrid operating model enables consultants, analysts, and engineers to work across distributed locations while relying heavily on cloud infrastructure, SaaS platforms, and collaborative DevSecOps pipelines. As the company expands into the European Union, it enters a more demanding regulatory and threat landscape shaped by the Digital Operational Resilience Act (DORA) and the NIS2 Directive[11].

These regulations significantly elevate expectations around cyber resilience, governance, and incident accountability. DORA requires organisations to demonstrate continuous operational resilience across ICT systems, cloud services, and third-party providers. NIS2 introduces stricter cybersecurity obligations, enhanced risk management requirements, mandatory incident reporting timelines, and direct accountability for senior leadership[11]. For an AI-driven organisation such as ResilientIntellect, these obligations directly affect the protection of intellectual property, AI training data, and customer trust.

To meet these challenges, this report proposes a strategic and forward-looking information security programme grounded in Zero Trust principles. Zero Trust rejects traditional perimeter-based security and instead assumes that no user, device, workload, or system is inherently trustworthy. This approach aligns naturally with ResilientIntellect's hybrid workforce, cloud-native infrastructure, and AI development pipelines. The programme is designed not only to achieve regulatory compliance but also to support business growth, operational agility, and innovation.

The proposed security model is structured around three Zero Trust boundaries: People, Pipelines, and Platforms. *People* include employees, consultants, and clients operating in a hybrid environment. *Pipelines* encompass AI models, data flows, and DevSecOps processes. *Platforms* represent cloud infrastructure, SaaS services, and third-party providers. This framing ensures that Zero Trust principles are applied consistently across all operational layers, creating a resilient and scalable security foundation for the organisation's EU expansion.

2. Organisational Structure: Embedding Zero Trust into People and Processes

A strong organisational structure is essential to ensure that Zero Trust principles are not only designed but also consistently implemented and governed[12]. For the new EU unit, ResilientIntellect requires a security structure that is lean, scalable, and tightly integrated with engineering, cloud operations, and compliance functions. Instead of treating security as a standalone technical function, the proposed structure embeds security leadership directly into operational and strategic decision-making, ensuring alignment with business goals and regulatory expectations[12].

At the top of the structure is the Security and Risk Steering Committee, reporting to the Board of Directors. This committee provides strategic oversight, ensures compliance with NIS2's governance and accountability requirements, and links cybersecurity decisions to business risk, operational resilience, and growth objectives[12]. Beneath this committee sits the Head of Information Security (EU), who defines the Zero Trust strategy, oversees regulatory compliance, and acts as the primary point of accountability for EU regulators, partners, and customers.

The Head of Information Security is supported by four specialist functions. The **Cloud Security Architect** designs and maintains the Zero Trust architecture across multi-cloud environments and SaaS platforms, ensuring consistent enforcement of identity-centric access controls, network segmentation, and continuous verification in line with DORA's ICT risk management requirements. The **DevSecOps Lead** secures AI development and deployment pipelines by embedding security controls into CI/CD workflows and managing software supply chain risks, including vulnerable dependencies and compromised AI models.

The **GRC and Compliance Lead** manages governance, risk, and compliance activities, including resilience testing, third-party risk assessments, and coordination of incident reporting within NIS2's mandated timelines. The **SOC and Incident Response Lead** is responsible for continuous monitoring, threat detection, and coordinated incident response, supporting operational resilience and ensuring incidents are identified, classified, and reported effectively.



Figure 3: Organisational Structure

Figure 3 illustrates this governance-led security structure, showing clear accountability from board-level oversight through the Security and Risk Steering Committee to the Head of Information Security (EU), supported by specialist roles in cloud security, DevSecOps, compliance, and incident response.

Overall, this structure ensures that security responsibilities are clearly defined, strategically governed, and aligned with ResilientIntellect’s AI-centric, cloud-native operating model. It provides a strong foundation for operationalising Zero Trust principles while supporting regulatory compliance and scalable business growth within the European Union.

3. Policies and Procedures: Translating Zero Trust into Everyday Practice

Zero Trust can only be effective if it is applied consistently in everyday work. For ResilientIntellect, policies and procedures translate high-level security principles into clear rules that guide how people use systems, develop AI, and access data. These policies are designed to protect sensitive assets while still allowing staff to work flexibly in a fast-moving, hybrid environment.

Rather than assuming that internal users, devices, or systems are trustworthy, all policies are built on the Zero Trust assumption that every access request could be a potential risk. Access is therefore continuously checked, limited to what is necessary, and removed when no longer required.

AI Usage Policy

The AI Usage Policy protects the confidentiality, integrity, and availability of ResilientIntellect's AI systems. All AI models must be trained using approved datasets with clearly documented sources. This reduces the risk of poisoned or low-quality data affecting threat intelligence outputs.

Training, testing, and production environments are kept separate to prevent unauthorised changes or accidental exposure of sensitive data. Every AI model is tracked throughout its lifecycle, from development to deployment, ensuring full accountability. Importantly, all AI-generated outputs are treated as untrusted until they are reviewed or validated. This reflects a Zero Trust mindset and helps reduce risks linked to manipulation, bias, or incorrect analysis[13].

Secure SaaS Configuration Policy

SaaS platforms are essential to ResilientIntellect's operations, but they also represent a major source of risk if poorly configured. This policy ensures that all SaaS services are configured securely by default and reviewed regularly.

Access to SaaS applications follows the principle of least privilege, meaning users only receive the permissions they need to do their job. Strong identity controls, including multi-factor authentication, are mandatory. Configuration changes are continuously monitored, and any insecure or unauthorised change is treated as a security incident. By assuming that SaaS platforms are not inherently secure, this policy aligns closely with Zero Trust principles and significantly reduces the likelihood of data exposure through misconfiguration[14].

BYOD Policy

The BYOD Policy supports ResilientIntellect's hybrid workforce while maintaining strong security controls. Personal devices are allowed to access corporate systems only if they meet defined security standards, such as up-to-date patches, encrypted storage, and active endpoint protection[15].

Access to internal resources is provided through Zero Trust Network Access (ZTNA) rather than traditional VPNs. This ensures that users are granted access only to specific applications, not the wider network, and only while their device remains compliant. Corporate data is isolated from personal data using secure containers, reducing the risk of data loss or leakage[15].

Policy	Key Risks Addressed	Zero Trust Enforcement
AI Usage Policy	• Poisoned models • Data leakage • Model manipulation	• Treat AI as untrusted • Dataset validation • Provenance tracking
Secure SaaS Configuration Policy	• Misconfigurations • Excess permissions • Account compromise	• Least privilege access • Continuous monitoring • Assume breach
BYOD Policy	• Untrusted endpoints • Data exfiltration • Malware infection	• Device posture checks • ZTNA (no VPN trust) • Continuous verification

Figure 4 Policy Table

Figure 4 illustrates how key organisational policies are mapped to specific risks and enforced through Zero Trust controls across AI systems, cloud services, and hybrid working environments.

4. Regulatory Compliance: Addressing DORA and NIS2 Requirements

The proposed Zero Trust security programme is designed to support compliance with both DORA and NIS2 as an integrated operational capability rather than a standalone compliance exercise. By embedding regulatory requirements directly into everyday security practices, the organisation strengthens resilience while reducing administrative overhead[16].

Under DORA, organisations must demonstrate digital operational resilience, including the ability to withstand, respond to, and recover from ICT-related incidents. Zero Trust directly supports these expectations by limiting lateral movement, reducing the blast radius of attacks, and enabling rapid isolation of compromised systems. Regular resilience testing, protected backups, and structured third-party risk assessments are built into operational workflows to meet DORA's requirements for ICT risk management, incident classification, and continuity planning[16].

NIS2 places strong emphasis on risk management, incident reporting, and executive accountability. The proposed governance structure ensures that cybersecurity risks are overseen at board level, meeting NIS2's expectations for leadership involvement. Continuous monitoring, threat detection, and coordinated incident response capabilities enable the organisation to identify and report significant incidents within mandated timelines[16]. Clearly defined roles and responsibilities ensure that regulatory notifications are accurate, timely, and aligned with legal obligations.

This integrated approach ensures that compliance is achieved naturally through strong security practices, while simultaneously improving the organisation's overall security maturity and operational resilience[16].

5. Governance and Strategy: Aligning Security with Business Goals Using NIST CSF 2.0

Effective cybersecurity governance is essential for ensuring that Zero Trust becomes a strategic, organisation-wide approach rather than a purely technical initiative. As ResilientIntellect expands into the European Union, governance must enable rapid business growth while meeting the heightened regulatory expectations introduced by DORA and the NIS2 Directive. To achieve this, the proposed security programme is aligned with the NIST Cybersecurity Framework (CSF) 2.0, which offers a structured, risk-based model for managing cybersecurity across modern, cloud-native environments[17].

NIST CSF 2.0 is particularly appropriate in this context because it elevates governance to a core function, mirroring NIS2's focus on senior management accountability[17]. By framing cybersecurity as a business risk rather than an isolated IT concern, the framework ensures that security considerations are embedded into strategic planning, investment decisions, and operational risk management.



Figure 5 : NIST Cybersecurity Framework (CSF) 2.0 core functions illustrating a lifecycle approach to governing, identifying, protecting, detecting, responding to, and recovering from cybersecurity risks.

As illustrated in figure 5, the **Govern** function establishes clear accountability, oversight, and decision-making structures[17]. Within ResilientIntellect's EU operations, governance is anchored by board-level oversight through the Security and Risk Steering Committee. This ensures that cybersecurity strategy remains aligned with business objectives and regulatory obligations. The Head of Information Security (EU) is

given explicit authority to enforce Zero Trust principles across people, pipelines, and platforms, directly supporting NIS2 governance requirements and DORA's ICT risk management expectations.

The **Identify** function focuses on understanding and prioritising risks to the organisation's most critical assets, including AI models, training data, threat intelligence outputs, cloud workloads, and SaaS platforms[17]. This includes identifying and mapping assets and data flows across internal systems, third-party providers, and software supply chains, reflecting DORA's focus on managing ICT third-party and concentration risk. By adopting a risk-informed perspective, the organisation can prioritise controls that protect high-value assets and reduce systemic vulnerabilities.

The **Protect** function translates governance decisions into practical Zero Trust controls[17]. Strong identity and access management, phishing-resistant MFA, device posture validation, and Zero Trust Network Access (ZTNA) safeguard hybrid workers. At the same time, least-privilege access and embedded security controls protect AI development environments and DevSecOps pipelines. These measures reduce the likelihood of credential compromise, unauthorised access, and lateral movement—without hindering operational agility.

Continuous monitoring across cloud infrastructure, SaaS platforms, endpoints, and AI workloads supports early detection of anomalous behaviour and potential security incidents[17]. Centralised logging and SIEM-driven analytics enable rapid incident identification and classification, directly supporting DORA's operational resilience requirements and NIS2's incident reporting obligations.

The **Respond** and **Recover** functions ensure that the organisation can manage incidents effectively and restore services quickly[17]. Clearly defined incident response procedures, supported by automation, enable swift containment and timely regulatory notification where required. Recovery activities prioritise the integrity and availability of AI assets through secure backups, resilience testing, and structured post-incident reviews. Insights gained from incidents are fed back into governance and risk management processes, strengthening long-term resilience.

By adopting NIST CSF 2.0 as the foundation for its governance model, ResilientIntellect embeds Zero Trust into strategic decision-making and organisational culture[17]. This approach enhances regulatory compliance, protects critical AI assets, and positions cybersecurity as a strategic enabler for secure growth and customer trust across the European Union.

6. Risk Management: Focus on AI and Cloud Threats

A high-level risk assessment highlights the most significant threats facing ResilientIntellect’s EU operations. Software supply-chain risks are particularly critical, as compromised open-source dependencies, malicious packages, or poisoned AI models could directly undermine the accuracy, reliability, and credibility of the organisation’s threat intelligence outputs[17]. Cloud misconfigurations represent another major risk, with the potential to expose sensitive data, disrupt services, or create unauthorised access paths across multi-cloud environments[18].

To mitigate these risks, the organisation should implement automated dependency scanning, maintain Software Bills of Materials (SBOMs), and enforce model validation and cryptographic signing to ensure the integrity of AI artefacts. Continuous cloud posture management and configuration monitoring further reduce the likelihood of misconfigurations going undetected[18]. Together, these controls reflect a proactive, Zero Trust-aligned approach to risk management, ensuring that threats are identified and addressed before they can impact operations.

Figure 6: High-Level Zero Trust–Aligned Risk Assessment for AI, Cloud, and Third-Party Assets

Risk Area	Threat Description	Likelihood	Impact	Risk Level	Key Mitigations (Zero Trust–Aligned)
AI Supply Chain	Poisoned training data or compromised AI models leading to inaccurate threat intelligence	Medium	High	High	Model validation, dataset provenance tracking, cryptographic signing of AI models
Software Dependencies	Vulnerable or malicious open-source libraries introduced via CI/CD pipelines	High	High	Critical	Automated dependency scanning, SBOMs, DevSecOps pipeline controls
Cloud Misconfiguration	Over-permissive IAM roles or exposed storage services	High	High	Critical	CSPM tools, least-privilege IAM, continuous configuration monitoring
SaaS Platform Risk	Misconfigured SaaS applications exposing sensitive client data	Medium	High	High	Secure SaaS configuration baselines, continuous monitoring, MFA enforcement
Third-Party Providers	Security weaknesses in cloud or SaaS vendors affecting service resilience	Medium	Medium	Medium	Third-party risk assessments, contractual security requirements, resilience testing

Figure 6

Figure 6 summarises the key cybersecurity risks facing ResilientIntellect’s EU operations, with a focus on AI systems, cloud infrastructure, SaaS platforms, and third-party providers. It assesses the likelihood and impact of each risk and highlights the Zero Trust–aligned controls used to reduce exposure and limit operational impact.

7. Contingency Planning: Ransomware Impacting AI Assets

In the event of a ransomware attack that encrypts AI training data or source code repositories, ResilientIntellect must respond rapidly to minimise operational disruption and reputational impact. Early detection is achieved through endpoint protection, behavioural analytics, and continuous monitoring. Once suspicious activity is identified, automated Zero Trust controls isolate affected devices, user accounts, and workloads to prevent lateral movement[19].

The incident response team then coordinates containment, eradication, and recovery activities. Recovery is supported by secure, off-network backups of AI datasets, model artefacts, and source code repositories. Before systems are restored, AI models and datasets undergo integrity validation to ensure they have not been tampered with or corrupted[19]. This step is essential for maintaining trust in the organisation's AI outputs.

Where required, regulatory notification under NIS2 is initiated, ensuring that significant incidents are reported within mandated timelines. After recovery, a structured post-incident review is conducted to identify root causes, evaluate the effectiveness of controls, and implement improvements across people, processes, and technology.

This contingency approach strengthens operational resilience, supports regulatory compliance, and ensures transparent and accountable incident handling.

8. Protection Mechanisms: Enforcing Zero Trust Through Modern Controls

To operationalise Zero Trust, ResilientIntellect must adopt a focused set of cloud-native security controls that enforce continuous verification and limit implicit trust. All access requests whether from users, devices, or workloads—are evaluated based on identity, device posture, and risk context before access is granted.

Cloud Security Posture Management (CSPM) provides continuous visibility across cloud environments, identifying misconfigurations, excessive permissions, and policy violations. This supports the Zero Trust principle of assuming breach by reducing attack surface and preventing configuration-based exposure.

Centralised monitoring is achieved through SIEM integrated with SOAR, enabling rapid detection, automated response, and containment of security incidents. This significantly reduces the time required to identify and respond to threats before they can spread across hybrid environments[20].

Endpoint Detection and Response (EDR) strengthen Zero Trust by continuously monitoring device behaviour and isolating compromised endpoints in real time. This is particularly important for preventing ransomware and lateral movement within the organisation[21].

Identity remains the core control point. Phishing-resistant multi-factor authentication (MFA) and Zero Trust Network Access (ZTNA) replace traditional VPN-based access, ensuring users are granted application-level access only when continuously verified[21].

Together, these controls provide a scalable and resilient Zero Trust security posture that supports business growth while meeting regulatory expectations under DORA and NIS2.

9. Conclusion

This report has outlined a Zero Trust-aligned security programme designed to support ResilientIntellect's expansion into the European Union. By embedding Zero Trust across people, pipelines, and platforms, the organisation adopts a security model that reflects the realities of hybrid working, cloud-native infrastructure, and AI-driven operations.

The proposed approach directly supports compliance with DORA and the NIS2 Directive by strengthening governance, improving operational resilience, and enabling rapid detection and response to cyber incidents. Board-level oversight and clearly defined accountability ensure that cybersecurity risk is treated as a strategic business concern rather than a purely technical issue.

In adopting this Zero Trust-aligned approach, ResilientIntellect is not simply responding to regulatory pressure, it is establishing cybersecurity as a core capability that underpins long-term growth, operational resilience, and competitive advantage in the European market.

References

- [1] D. Shannon. "The Acceleration Of AI: Keep Up Or Be Left Behind." Accessed: Dec 22 2025. [Online]. Available: <https://www.forbes.com/councils/forbestechcouncil/2025/04/03/the-acceleration-of-ai-keep-up-or-be-left-behind/>
- [2] P. Ashok and C. Das, "Artificial Intelligence Powered Product Management: Leveraging Data-Driven Insights for Growth, Retention, and Monetization", pp. 1322–1327. [Online].doi: 10.1109/ICoEIT63558.2025.11211726.
- [3] D. Fernández Llorca, V. Charisi, R. Hamon, I. Sánchez and E. Gómez, "Liability Regimes in the Age of AI: a Use-Case Driven Analysis of the Burden of Proof," vol. 76, pp.613–644, 2023, doi: 10.1613/jair.1.14565.
- [4] Unknown, "What is AI Governance? | IBM," -10-10 2024, Accessed: Dec 22, 2025. Available at: <https://www.ibm.com/think/topics/ai-governance>.
- [5] B. C. Cheong, "Transparency and accountability in AI systems: safeguarding wellbeing in the age of algorithmic decision-making," *Front. Hum. Dyn.* vol. 6, July 3 2024, doi: 10.3389/fhumd.2024.1421273. Accessed: Dec 27, 2025. Available at: <https://www.frontiersin.org/journals/human-dynamics/articles/10.3389/fhumd.2024.1421273/full>.
- [6] S. Vengathattil and S. M. Shaffi, "Ethical Implications of AI-Powered Decision Support Systems in Organizations", pp. 105–112. [Online].doi: 10.1109/ICAIDE65466.2025.11189693.
- [7] I. Atoum, "Revolutionizing AI Governance: Addressing Bias and Ensuring Accountability Through the Holistic AI Governance Framework," vol. 16, no. 2, 2025, doi: 10.14569/IJACSA.2025.0160283.
- [8] Boskir E , Chen Y, Ghulam S , Ma Q, Razmi R, Swaminathan R, Thiru I, "Predictive Policing or Predictive Prejudice? A Study of the Legal, Historical and Ethical Implications of AI in Policing | OxJournal," 18 Aug 2025, Accessed: Dec 27, 2025. Available at: <https://www.oxjournal.org/predictive-policing-or-predictive-prejudice/>.
- [9] A. Omran Almagrabi and R. A. Khan, "Optimizing Secure AI Lifecycle Model Management With Innovative Generative AI Strategies," *Access*. vol. 13, pp.12889–12920, 2025, doi: 10.1109/ACCESS.2024.3491373.
- [10] W. Villegas-Ch and J. García-Ortiz, "Toward a Comprehensive Framework for Ensuring Security and Privacy in Artificial Intelligence," vol. 12, no. 18, pp.3786, 2023, doi: 10.3390/electronics12183786.

- [11] F. Teichmann, "Cybersecurity of critical infrastructure in europe: the NIS2 directive in focus," *Int.Cybersecur.Law Rev.* vol. 6, no. 3, pp.207–220, 2025, doi: 10.1365/s43439-025-00154-4.
- [12] Y. Bobbert and J. Scheerder, "Zero Trust Validation: from Practice to Theory : An empirical research project to improve Zero Trust implementations", pp. 93–104. [Online].doi: 10.1109/STC55697.2022.00021.
- [13] unknown, "Artificial Intelligence Usage Policy," Accessed: Jan 2, 2026. Available at: <https://www.freemanclarke.com/en-gb/wp-content/uploads/sites/1/2023/10/Artificial-Intelligence-Usage-Policy.pdf>.
- [14] unknown, "Create SaaS Policy Rule Recommendations," Accessed: Jan 4, 2026. Available at: <https://docs.paloaltonetworks.com/saas-security/saas-security-inline/manage-saas-security-inline-policy/create-saas-policy-rule-recommendations>.
- [15] unknown, "Bring Your Own Device (BYOD) Policy Best Practices | Okta," Accessed: Jan 3, 2026. Available at: <https://www.okta.com/identity-101/byod-policy/>.
- [16] K. Mentzelou, P. T. Chountalas, F. C. Kitsios, A. I. Magoutas and T. K. Dasaklis, "Identifying and Modeling Barriers to Compliance with the NIS2 Directive: A DEMATEL Approach," vol. 5, no. 4, pp.97, 2025, doi: 10.3390/jcp5040097.
- [17] unknown, "The NIST Cybersecurity Framework (CSF) 2.0," 26 February 2024, Accessed: Jan 2, 2026. Available at: <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf>.
- [18] R. Kummara, "Cloud Security Using AI: Transforming Digital Protection In The Modern Era," vol. 8, pp.27–38, 2025, doi: 10.63278/jicrcr.vi.3215.
- [19] G. Baran. "AI-Powered Ransomware Is the Emerging Threat That Could Bring Down Your Organization." Accessed: Jan 2 2026. [Online]. Available: <https://cybersecuritynews.com/ai-powered-ransomware/>
- [20] A. Hassan, A. Rauf, N. Shafqat, R. Latif and H. Khan, "ZenGuard a machine learning based zero trust framework for context aware threat mitigation using SIEM SOAR and UEBA," *Sci Rep.* vol. 15, no. 1, pp.35871, October 14 2025, doi: 10.1038/s41598-025-20998-4. Accessed: Jan 2, 2026. Available at: <https://www.nature.com/articles/s41598-025-20998-4>.
- [21] O. S. "Zero Trust Security: Implementation Guide for Modern Businesses (2025)." Accessed: Jan 4 2026. [Online]. Available: <https://gettrusted.io/blog/zero-trust-security-implementation-guide-for-modern-businesses-2025/>

FINAL GRADE

GENERAL COMMENTS

87 /100

PAGE 1

PAGE 2

PAGE 3

PAGE 4

PAGE 5

PAGE 6

PAGE 7

PAGE 8

PAGE 9

PAGE 10

PAGE 11

PAGE 12

PAGE 13

PAGE 14

PAGE 15

PAGE 16

GRADING FORM: L5 DSC (ACCA5034)

REHAN JOSEPH

87

ENGAGEMENT WITH LITERATURE

All work must be appropriately researched and referenced with an extensive range of relevant and credible literature. High-level referencing skills consistently applied. Most references should be current and Laws must cite current legislation in the appropriate government site(s). (20 marks)



Excellent engagement with academic and authoritative sources, well integrated into the discussion. Literature is relevant and generally synthesised effectively.

18

KNOWLEDGE AND UNDERSTANDING

You will need to show application of skills and research in the assignment by demonstrating critical understanding of the well-

established concepts and principles of the subject and the main methods of enquiry. May go beyond established theories.(10 marks)



Demonstrates excellent conceptual understanding of AI governance, ethics, lifecycle risk, and compliance. Terminology is accurate and explanations are clear and confident.

9

COGNITIVE AND INTELLECTUAL

Marks will be awarded: For critical analysis and evaluation, based on the quality of the documentation/report/paper. How relevant conclusions are drawn and justified. (25 marks)



Sustained and well-structured critical analysis across Task 1 and Task 2. Clear justification of governance choices, strong ethical reasoning, and effective linkage between risk, compliance, and strategy. Analysis is consistent, though still cautious in places.
A genuinely high-quality script. Strong across all criteria, with clear artefact evidence and sustained critical depth.

22

PRACTICAL SKILLS

Marks will be awarded based on the quality of and level of detail provided for the design approach; diagrams, tables, charts and any other design methods utilise will enhance marks. These should be properly explained and contextualised. (25 marks)



Clear organisation structure diagram
Multiple figures, tables, and frameworks
Diagrams are well labelled and integrated

22

TRANSFERABLE SKILLS FOR LIFE

This will be assessed based on the written and presentational quality of your report: • Flow • Professionalism • Comprehension (grammar and spelling) • Completeness and presentation of report (20 marks)



Well written, clearly structured, and academically presented. Paragraph justification is inconsistent, but overall professionalism and clarity are very good.

16