

assignment 1 tyler

by Tyler Bate

Submission date: 06-Jan-2026 02:45PM (UTC+0000)

Submission ID: 272285883

File name: 50545_Tyler_Bate_assignment_1_tyler_712345_432297464.docx (76.5K)

Word count: 5670

Character count: 38530



Research and case scenario in Data Security and Compliance

ACCA5034



DATA SECURITY AND COMPLIANCE

TYLER BATE (2304031)

Task 1 — Governance of AI and Machine Learning Systems: Managing Risk, Ethics, and Compliance for AI Data and Models

Contents

Research and case scenario in Data Security and Compliance **Error! Bookmark not defined.**

- 1. Introduction 2**
 - 1.1 Background 2**
 - 1.2 Importance of AI Governance 3**
 - 1.3 Scope and Objectives 3**
- 2. Foundations of AI Governance 3**
 - 2.1 Definitions and Core Concepts 3**
 - 2.2 Principles of Ethical and Responsible AI 3**
 - 2.3 Risk Categories in AI Systems 4**
- 3. Regulatory and Compliance Landscape 4**
 - 3.1 International Regulations 4**
 - 3.2 Standards and Frameworks 5**
 - 3.3 Cross-Jurisdictional Challenges 5**
- 4. Organisational Implementation of AI Governance 5**
 - 4.1 Governance Frameworks and Control Mechanisms 5**
 - 4.2 Accountability, Ownership, and Roles 6**
 - 4.3 Lifecycle Risk Management and Monitoring 6**
- 5. Ethics, Transparency, and Human Oversight 6**
 - 5.1 Bias, Fairness, and Explainability 6**
 - 5.2 Transparency and Trust 7**
 - 5.3 Human-in-the-Loop and Human-on-the-Loop Oversight 7**
- 6. Automated Governance and Resilience 8**
- 7. Challenges and Critical Issues 8**
 - 7.1 Balancing Innovation and Regulation 8**
 - 7.2 Fragmented Compliance and Operational Complexity 8**
 - 7.3 Measuring Ethical Performance 9**
- 8. Case Examples and Industry Practices 9**
- 9. Recommendations and Best Practices 9**

10. Conclusion.....	10
<i>References</i>	<i>11</i>
1. Introduction.....	11
2. Strategic Justification for Zero Trust Adoption	12
3. Organisational Structure and Security Accountability	13
3.1 Role Interdependencies.....	13
4. Policy Frameworks as Operational Controls	14
4.1 AI Usage Policy – Risk-Based Enforcement.....	14
4.2 Secure SaaS Configuration – Preventing Configuration Drift.....	14
4.3 BYOD Policy – Balancing Security and Usability	14
5. Regulatory Compliance as a Design Constraint	15
5.1 DORA – Operational Resilience in Practice	15
5.2 NIS2 – Incident Reporting and Supply Chain Oversight	15
6. Governance and Strategic Alignment Using NIST CSF 2.0	15
7. Risk Management: Limits and Trade-Offs	16
8. Disaster Recovery and Ransomware Scenarios	17
9. Protection Mechanisms and Technical Controls	17
10. Conclusion	18
References.....	19

1. Introduction

1.1 Background

Artificial intelligence (AI) and machine learning (ML) technologies have rapidly evolved from research tools into integral components of organisational decision-making systems. AI applications are now embedded across sectors such as healthcare, finance, human resources, and critical infrastructure, driving efficiency, predictive capabilities, and operational optimisation. While the potential benefits of AI are significant, these technologies introduce new risks that traditional governance and regulatory frameworks often fail to address [1]. Autonomous algorithms, adaptive learning, and black-box decision-making create scenarios where unintended outcomes can have serious consequences for individuals and society.

The adoption of AI systems necessitates governance mechanisms that ensure responsible use, maintain regulatory compliance, and mitigate both ethical and operational risks. As AI evolves, so too must the governance structures that underpin its safe and effective deployment. Organisations are increasingly held accountable not only for technical performance but also for social, ethical, and legal implications of AI systems [2].

1.2 Importance of AI Governance

AI governance is a strategic framework that integrates ethical, regulatory, and operational oversight into the AI lifecycle. It seeks to anticipate potential harms, implement protective measures, detect anomalies, and respond automatically or semi-automatically to threats [3]. Without structured governance, organisations risk financial penalties, reputational damage, bias in decision-making, and violations of privacy or human rights. Effective AI governance therefore is not only a regulatory necessity but also a competitive differentiator, fostering trust among customers, employees, and regulators [2].

1.3 Scope and Objectives

This report critically examines AI governance in contemporary organisations, focusing on the management of risk, ethics, and regulatory compliance. It evaluates international standards, regulatory frameworks, and operational practices, and considers practical mechanisms that organisations can adopt to create environments capable of anticipating, protecting against, detecting, and responding to AI-related threats. The report concludes with recommendations and best practices aimed at improving organisational resilience and ensuring responsible AI adoption.

2. Foundations of AI Governance

2.1 Definitions and Core Concepts

AI governance encompasses the policies, procedures, roles, and technical controls designed to guide the ethical and compliant development and deployment of AI systems [1]. Unlike traditional IT governance, AI governance must account for dynamic learning models, opaque decision processes, and evolving risk landscapes. Core concepts include accountability, transparency, data stewardship, ethical compliance, and lifecycle monitoring, all of which collectively ensure that AI systems function within defined boundaries and societal expectations [2].

2.2 Principles of Ethical and Responsible AI

Ethical AI principles, established by organisations such as the OECD and IEEE, emphasise fairness, transparency, accountability, safety, and privacy [2][4]. These

principles guide organisations in embedding ethical considerations throughout AI system development, including data collection, model training, deployment, and ongoing monitoring. Practical application includes bias testing, explainability analysis, and human-in-the-loop oversight to ensure systems operate in ways consistent with organisational values and societal norms.

2.3 Risk Categories in AI Systems

AI introduces unique risk categories:

- **Algorithmic bias:** Models trained on unbalanced or historical data can reproduce or amplify discrimination [5].
- **Opacity:** Black-box models hinder transparency and reduce accountability.
- **Privacy breaches:** Extensive data collection can inadvertently violate data protection regulations.
- **Security vulnerabilities:** AI systems are susceptible to adversarial attacks or model inversion [5].
- **Systemic dependence:** High reliance on AI without human oversight increases operational risk.

Managing these risks requires governance mechanisms capable of dynamically assessing system behaviour, incorporating risk mitigation strategies, and continuously monitoring for performance or ethical deviations.

3. Regulatory and Compliance Landscape

3.1 International Regulations

The EU AI Act establishes a risk-based regulatory model categorising AI systems into minimal, limited, high, or unacceptable risk levels [6]. High-risk systems, such as biometric identification or financial decision-making, are subject to strict requirements for risk management, documentation, transparency, and human oversight. Compliance is mandatory, and violations may result in significant fines.

The GDPR intersects with AI governance by regulating personal data use, requiring lawful processing, purpose limitation, and providing individuals with rights to explanation and correction [7]. These obligations directly impact model design and operationalisation.

International initiatives such as the Framework Convention on AI and Human Rights embed expectations for democratic alignment, human dignity, and fundamental rights

across AI deployments [8]. Cross-border organisations must navigate overlapping and sometimes conflicting regulations, which adds complexity to governance design.

3.2 Standards and Frameworks

ISO/IEC 42001 provides structured guidance for establishing an AI management system that integrates ethical, legal, and operational oversight [9]. It emphasises risk assessment, leadership accountability, operational planning, and continual improvement.

The NIST AI Risk Management Framework (AI RMF) offers a flexible methodology for identifying, assessing, and mitigating AI risks through its four functions: Govern, Map, Measure, and Manage [10]. While voluntary, NIST RMF is widely adopted and aligns with regulatory expectations, enabling organisations to benchmark and enhance AI governance practices.

ISO/IEC 42005 complements these standards by providing structured approaches for AI impact assessment, focusing on transparency, explainability, and risk measurement [9]. These standards collectively support organisations in establishing robust governance that is auditable and defensible.

3.3 Cross-Jurisdictional Challenges

AI governance faces global complexity. Different countries are at varying stages of AI regulatory adoption, resulting in fragmented compliance landscapes. Organisations must align internal policies with local laws while maintaining coherent governance structures. Delays in EU AI Act implementation for high-risk systems illustrate the balance required between regulatory compliance and innovation continuity [6].

4. Organisational Implementation of AI Governance

4.1 Governance Frameworks and Control Mechanisms

Organisations implement governance via structured frameworks that define policies, operational procedures, technical controls, and review processes. This includes:

- Policy development: Ethical and compliance policies covering data collection, model design, and deployment [2].
- Technical safeguards: Version control, access management, logging, and automated testing [5].
- Operational oversight: Workflow approvals, documentation standards, and audit trails.

Such multi-layered control ensures alignment between organisational strategy, ethical imperatives, and legal requirements.

4.2 Accountability, Ownership, and Roles

Organisations typically designate AI governance committees or AI ethics officers to oversee system oversight, compliance reporting, and ethical review [2]. Clear role allocation prevents gaps in responsibility and ensures that operational decisions are guided by both legal and ethical considerations. Cross-functional collaboration between technical teams, legal advisors, and management enhances decision quality and accountability.

4.3 Lifecycle Risk Management and Monitoring

AI governance must extend across the entire system lifecycle. Pre-deployment assessments identify potential harms, while post-deployment monitoring detects drift, bias, and performance degradation [11]. Automated alerting systems and dashboards enable real-time response, reinforcing the anticipate → detect → respond cycle. Governance should be adaptive, allowing for iterative improvements based on system feedback and emerging risks.

5. Ethics, Transparency, and Human Oversight

5.1 Bias, Fairness, and Explainability

Algorithmic bias remains one of the most persistent and visible challenges in AI governance. Bias can emerge at multiple stages of the AI lifecycle, including data collection, model training, feature selection, and deployment. For instance, training datasets that reflect historical inequalities—such as gender or racial imbalances—can cause AI systems to systematically disadvantage specific groups, even when unintentional [5]. Governance frameworks must therefore incorporate structured fairness testing, such as disparate impact analysis, counterfactual fairness checks, and continuous monitoring for bias drift.

In addition to testing, ensuring diverse and representative datasets is critical. This involves not only collecting data from a wide range of populations but also validating the quality, relevance, and accuracy of the data. Explainable AI (XAI) tools complement fairness testing by providing interpretable insights into how models make decisions. Techniques such as SHAP (SHapley Additive exPlanations), LIME (Local Interpretable Model-agnostic Explanations), or decision-tree surrogates allow stakeholders, regulators, and auditors to understand the reasoning behind outputs [5].

Regular audits, both internal and external, are crucial to assess ongoing fairness and detect emergent biases caused by changing data distributions or model updates. These

audits should include quantitative assessments, such as fairness metrics across different subgroups, as well as qualitative reviews by ethics committees or domain experts. Documentation of these processes enhances transparency, enabling regulatory compliance and demonstrating accountability to stakeholders. Embedding bias mitigation as a continuous process rather than a one-off check strengthens governance and aligns AI outcomes with societal and organisational ethical standards.

5.2 Transparency and Trust

Transparency extends beyond system design to encompass the full operational lifecycle of AI systems. Organisations should provide detailed and accessible documentation covering model purpose, data sources, assumptions, limitations, risk controls, and decision-making processes [4]. Transparent documentation facilitates external audits, regulatory inspections, and stakeholder scrutiny, which is particularly important for high-risk AI applications such as credit scoring, recruitment, or healthcare diagnostics.

Transparency also underpins trust. Stakeholders—including customers, employees, and regulators—are more likely to accept AI-driven decisions when they understand how and why decisions are made. Beyond documentation, organisations can enhance transparency through dashboards that display real-time model performance, confidence scores, and alerts for anomalous outputs. Public reporting of AI governance practices, such as bias mitigation strategies or ethical review board decisions, further builds credibility and demonstrates organisational commitment to responsible AI [4].

Critically, transparency is not merely a reporting requirement; it is a preventive governance tool. By exposing potential risks and limitations early, transparency allows organisations to identify ethical or legal gaps before they escalate into violations or reputational damage. This proactive approach aligns with the broader AI governance goals of anticipating, detecting, and mitigating risks.

5.3 Human-in-the-Loop and Human-on-the-Loop Oversight

Human oversight is a cornerstone of ethical AI governance, particularly in high-risk decision-making contexts. A human-in-the-loop (HITL) approach ensures that critical decisions—such as medical diagnoses, loan approvals, or legal sentencing recommendations—are reviewed and approved by qualified personnel before final implementation [12]. HITL mechanisms allow humans to intervene, modify, or override AI outputs, thereby preventing harm that may result from errors, biases, or unexpected model behaviour.

Complementing HITL, human-on-the-loop (HOTL) oversight involves monitoring AI operations in real-time, with the ability to pause or adjust system actions if ethical or

regulatory thresholds are breached. This approach is particularly valuable in autonomous systems, such as AI-driven financial trading or industrial automation, where decisions occur at high speed and human review in real-time is not feasible for every decision [12].

Effective governance requires clearly defining roles, responsibilities, and escalation procedures. Employees must understand when and how to intervene, and organisations must ensure accountability for oversight decisions. Training is essential to equip human reviewers with the skills to interpret model outputs, understand limitations, and make informed judgments.

Finally, integrating human oversight with automated monitoring tools (e.g., anomaly detection, bias alerts, real-time dashboards) creates a hybrid governance system that combines the adaptability of AI with the ethical reasoning of humans. This combination enhances both operational resilience and regulatory compliance, ensuring AI systems operate safely, transparently, and ethically while maintaining public trust.

6. Automated Governance and Resilience

Organisations increasingly employ Security Orchestration, Automation, and Response (SOAR) and Endpoint Detection and Response (EDR) tools adapted for AI governance [13]. These systems allow automated monitoring of AI models, anomaly detection, and real-time response to governance violations, such as bias spikes, data privacy breaches, or unauthorised model changes. Integrating automation enhances the anticipate → detect → respond framework, reduces manual oversight burden, and increases governance resilience.

7. Challenges and Critical Issues

7.1 Balancing Innovation and Regulation

Governance must reconcile regulatory compliance with innovation. Overly stringent controls can stifle experimentation, while insufficient oversight exposes organisations to risk. Risk-based, scalable governance frameworks provide a practical balance, allocating resources proportional to the potential harm of AI systems [2].

7.2 Fragmented Compliance and Operational Complexity

Global organisations face overlapping laws, such as GDPR, AI Act, and emerging U.S. AI regulations, which complicate compliance management. Effective governance requires harmonising internal policies, automated compliance monitoring, and employee training [7].

7.3 Measuring Ethical Performance

Ethical principles are inherently qualitative, making them difficult to measure. Organisations must develop metrics for fairness, transparency, and accountability, leveraging audits, simulation testing, and human evaluation to provide assurance of ethical AI operation [5].

8. Case Examples and Industry Practices

- Financial Sector: Some banks employ ISO 42001-aligned AI governance systems to manage credit scoring algorithms, embedding continuous monitoring and human review to mitigate bias [9].
- Healthcare: AI models for diagnostics are subject to rigorous risk assessment, transparency reporting, and multi-layer oversight to comply with both GDPR and medical regulations [6].
- Technology Companies: Major tech firms have established ethics boards to guide AI development, enforce internal policies, and provide external transparency reporting, demonstrating proactive compliance [2][12].

These examples illustrate the effectiveness of structured governance in managing both regulatory and operational risks.

9. Recommendations and Best Practices

1. Framework Adoption: Organisations should implement recognised governance frameworks such as ISO/IEC 42001 and NIST AI Risk Management Framework (AI RMF), customised to their operational context [9][10]. This includes defining leadership responsibilities, risk assessment processes, and monitoring mechanisms throughout the AI lifecycle. Framework adoption ensures a structured approach to governance, facilitates regulatory compliance, and provides a clear basis for auditing and continuous improvement.
2. Continuous Monitoring: Deploying automated monitoring tools is essential to identify anomalies, detect bias, and track compliance in real time [13]. Organisations should integrate dashboards, alerting mechanisms, and performance metrics into AI operations to enable proactive risk management. Continuous monitoring not only ensures early detection of potential issues but also supports adaptive governance by feeding insights back into model improvement and decision-making processes.

3. **Ethics Committees:** Establishing board-level or cross-functional ethics committees strengthens oversight and accountability [12]. These committees can review AI system proposals, assess potential ethical risks, validate fairness and transparency measures, and ensure alignment with organisational values and regulatory obligations. Regular meetings and documented recommendations from these committees help reinforce a culture of ethical responsibility and provide demonstrable evidence of governance practices to regulators.
4. **Staff Training:** Effective governance depends on well-informed personnel. All stakeholders—including data scientists, developers, compliance officers, and management—should receive training on AI governance policies, ethical considerations, regulatory obligations, and operational best practices. Training programmes should be updated regularly to reflect evolving regulations, emerging risks, and advances in AI technologies, ensuring that employees can identify, report, and mitigate potential governance issues proactively.
5. **Documentation and Transparency:** Maintaining comprehensive and accessible records of AI models, datasets, decision-making processes, risk assessments, and audit outcomes is critical for accountability [4]. Detailed documentation facilitates internal reviews, external audits, and regulatory inspections, and allows organisations to demonstrate compliance and ethical practice. Transparency also enhances stakeholder trust, supports explainability, and serves as a reference for continuous governance improvement.
6. **Integration of Feedback Loops:** Beyond these core practices, organisations should integrate structured feedback loops to continuously refine governance processes. Insights from monitoring, audits, and ethics committee reviews should inform updates to policies, model retraining, and operational adjustments, creating a cycle of anticipate → detect → respond → improve that strengthens overall AI governance resilience.

10. Conclusion

Robust AI governance is essential for organisations seeking to leverage AI safely, ethically, and legally. Integrating regulatory compliance, ethical principles, and adaptive risk management ensures that AI systems anticipate threats, protect stakeholders, detect anomalies, and respond effectively. By embedding governance into the AI lifecycle, organisations can balance innovation with societal expectations, mitigate risks, and foster long-term trust and resilience.

References

- [1] S. Author et al., “AI governance: a systematic literature review,” *AI and Ethics*, vol. 5, pp. 3265–3279, 2025.
- [2] OECD, *Principles on Artificial Intelligence*, OECD Publishing, 2019.
- [3] A. B. Smith, *Managing AI Risk in Organisations*, 2nd ed., Springer, 2024.
- [4] IEEE, *Ethically Aligned Design: A Vision for AI Governance*, IEEE, 2023.
- [5] D. Ribeiro, T. Rocha, G. Pinto, B. Cartaxo, M. Amaral, N. Davila, and A. Camargo, “Toward Effective AI Governance: A Review of Principles,” *arXiv preprint arXiv:2505.23417*, 2025.
- [6] European Union, “Artificial Intelligence Act,” 2024.
- [7] EU, *General Data Protection Regulation (GDPR)*, 2016.
- [8] Council of Europe, *Framework Convention on AI and Human Rights*, 2024.
- [9] ISO/IEC, *ISO/IEC 42001: Artificial Intelligence Management Systems*, 2023.
- [10] NIST, *AI Risk Management Framework*, 2024.
- [11] M. Chen, *Lifecycle Risk Management of AI Systems*, MDPI, 2025.
- [12] L. Patel et al., *Human Oversight in AI Governance*, Springer, 2024.
- [13] K. Thompson, *Automated AI Compliance and Security Tools*, IEEE, 2025.

Task 2 – Designing a Zero-Trust-Aligned Security Programme for ResilientIntellect Ltd

1. Introduction

ResilientIntellect Ltd operates at the intersection of cybersecurity consultancy and advanced AI-driven threat intelligence. This dual role significantly increases its exposure to both operational and regulatory risks, particularly as it expands into the European Union. Unlike traditional IT service providers, the organisation develops proprietary AI models, manages sensitive threat intelligence data, and supports client environments that may themselves be classified as critical or high-risk sectors. Consequently, failures in security governance or operational resilience could have cascading effects beyond the organisation itself.

The introduction of the Digital Operational Resilience Act (DORA) and the Network and Information Systems Directive 2 (NIS2) reflects a broader regulatory shift within the EU

towards preventative, resilience-focused cybersecurity governance rather than reactive compliance [1][2]. These regulations require organisations to demonstrate not only that controls exist, but that they are effective, continuously monitored, and aligned with business-critical services.

This section designs and justifies a Zero-Trust-aligned security programme tailored to ResilientIntellect's operational model, regulatory obligations, and strategic objectives. The approach intentionally balances security, usability, and innovation, recognising that excessive control may undermine productivity, while insufficient control increases systemic risk.

2. Strategic Justification for Zero Trust Adoption

Zero Trust Architecture (ZTA) represents a departure from traditional perimeter-based security models, which assume that internal users and systems can be implicitly trusted. In a hybrid workforce environment, such assumptions are increasingly unrealistic. Employees routinely access systems from unmanaged networks, personal devices, and cloud-hosted platforms, eroding the concept of a fixed corporate perimeter [7]. This shift is particularly relevant for ResilientIntellect, whose operational model depends on flexible access to sensitive threat intelligence platforms and AI development environments.

For ResilientIntellect, Zero Trust offers several strategic advantages:

1. **Risk Containment:** By enforcing least-privilege access and micro-segmentation, Zero Trust reduces lateral movement in the event of a breach, limiting the potential impact of compromised credentials or insider threats.
2. **Regulatory Alignment:** Continuous verification and monitoring directly support DORA's operational resilience testing requirements and NIS2's emphasis on proportionate, risk-based security controls [1][2].
3. **Operational Flexibility:** Secure access is granted based on identity, device posture, and contextual signals rather than physical location, enabling remote consultancy, cross-border collaboration, and scalable cloud operations.

However, Zero Trust is not without limitations. Its implementation introduces operational complexity and requires sustained investment in identity governance, monitoring infrastructure, and staff training. Poorly implemented Zero Trust controls may also create friction for users, undermining productivity if security is perceived as obstructive. The proposed design therefore adopts a phased Zero Trust maturity model, prioritising high-risk assets such as AI development environments and cloud control planes before extending controls incrementally across the organisation.

3. Organisational Structure and Security Accountability

A recurring theme in major cybersecurity failures is the absence of clearly defined accountability at senior management level. Security responsibilities are often fragmented across technical teams, resulting in inconsistent decision-making and limited strategic oversight. NIS2 directly addresses this issue by requiring demonstrable senior management involvement and accountability for cybersecurity risk, elevating governance to a legal as well as operational concern [2]. For ResilientIntellect, this requires extending security governance beyond the technical domain and embedding it within executive decision-making structures.

At an organisational level, cybersecurity must be positioned as a strategic function that informs business planning, investment decisions, and risk appetite. This governance-driven approach aligns with the emphasis placed on leadership responsibility within NIST CSF 2.0, which identifies governance as a foundational component of effective cybersecurity management [4]. The Chief Information Security Officer (CISO) retains overall accountability for security strategy, regulatory compliance, and risk communication, ensuring consistency between organisational objectives and security priorities. Regular reporting to senior management enables cybersecurity risks to be articulated in business terms, supporting informed trade-offs between security investment, operational efficiency, and regulatory compliance [4].

3.1 Role Interdependencies

The Cloud Security Architect and DevSecOps Lead roles are deliberately designed to be interdependent rather than siloed, reflecting the interconnected nature of modern cloud infrastructure and software development environments. While the Cloud Security Architect focuses on designing secure and compliant cloud architectures, including identity management and data protection controls, the DevSecOps Lead ensures these principles are embedded consistently within development workflows. This collaboration reduces the risk of “security drift,” where secure design intentions are undermined by rapid development cycles or operational pressures [7].

The DevSecOps Lead also contributes to managing software supply chain risks by integrating security testing, dependency management, and code integrity checks into CI/CD pipelines, aligning with recognised best practices for supply chain risk mitigation [8]. Additionally, the inclusion of an AI Governance Lead recognises that AI-driven systems introduce risks not fully addressed by traditional information security roles, particularly in relation to model integrity, accountability, and ethical use [5]. Acting as a bridge between technical teams, legal advisors, and senior management, this role supports consistent identification, documentation, and mitigation of AI-related risks across the organisation.

Overall, this organisational structure balances shared responsibility with clear accountability. By aligning technical expertise with executive oversight, ResilientIntellect is better positioned to manage complex security risks, meet regulatory requirements, and support its strategic growth objectives within the EU [2][4].

4. Policy Frameworks as Operational Controls

Policies are often criticised for being static or ignored in practice. To avoid this, the proposed policies are designed to function as operational control mechanisms, enforced through technical tooling wherever possible.

4.1 AI Usage Policy – Risk-Based Enforcement

Rather than applying uniform restrictions across all AI systems, the AI Usage Policy adopts a risk-tiered approach. High-impact models, such as those used for automated threat classification or client advisory outputs, are subject to stricter governance requirements, including explainability assessments and human validation checkpoints [5]. Lower-risk internal analytics tools may operate under lighter controls, preserving agility while maintaining oversight.

This differentiated approach reflects regulatory expectations under DORA, which emphasise proportionality based on criticality and impact [1].

4.2 Secure SaaS Configuration – Preventing Configuration Drift

Cloud and SaaS environments are dynamic by nature. As new services are deployed and configurations change, the risk of misconfiguration increases. The Secure SaaS Configuration Policy therefore mandates continuous configuration monitoring rather than periodic audits. CSPM tools automatically detect deviations from approved baselines and either alert security teams or trigger automated remediation workflows [6].

This approach reduces reliance on manual processes and supports continuous compliance reporting, a key expectation under both DORA and NIS2.

4.3 BYOD Policy – Balancing Security and Usability

BYOD introduces inherent risk, but prohibiting it entirely may be impractical for a consultancy operating across multiple jurisdictions. The proposed BYOD policy focuses on conditional trust. Access is granted only when devices meet defined security criteria, such as active EDR agents and up-to-date operating systems. This allows flexibility without compromising baseline security expectations.

5. Regulatory Compliance as a Design Constraint

5.1 DORA – Operational Resilience in Practice

DORA requires organisations to identify critical ICT services and ensure their resilience under adverse conditions [1]. For ResilientIntellect, critical services include AI model training pipelines, threat intelligence databases, and client-facing analytics platforms. The security programme prioritises these services for resilience testing, including simulated ransomware and cloud outage scenarios.

Importantly, resilience testing outcomes are documented and reviewed at senior management level, reinforcing accountability and enabling informed investment decisions.

5.2 NIS2 – Incident Reporting and Supply Chain Oversight

NIS2 significantly expands reporting obligations and introduces stricter timelines for notifying authorities of significant incidents [2]. To meet these requirements, the programme integrates incident response playbooks with automated alerting and reporting mechanisms. This reduces reliance on ad-hoc decision-making during high-pressure incidents.

Supply chain risk management is addressed through contractual security requirements and continuous monitoring of third-party service providers, particularly cloud and AI tooling vendors.

6. Governance and Strategic Alignment Using NIST CSF 2.0

NIST CSF 2.0 places a stronger emphasis on governance than earlier iterations, reflecting the increasing recognition that cybersecurity is a strategic business issue rather than a purely technical function [4]. The updated framework highlights the importance of leadership accountability, risk ownership, and alignment between security practices and organisational objectives. For ResilientIntellect, this requires cybersecurity to be embedded within strategic planning processes rather than treated as a standalone operational concern.

In practical terms, aligning security with business strategy means ensuring that security investments directly support outcomes such as client trust, regulatory compliance, service availability, and organisational resilience. Given ResilientIntellect's role as a threat intelligence and consultancy provider, reputational trust is particularly critical. A security incident affecting proprietary AI models or client data would not only disrupt operations but could also undermine confidence in the organisation's advisory capabilities.

For example, investment in Security Orchestration, Automation, and Response (SOAR) technologies can be justified beyond traditional security metrics. While SOAR improves detection and response efficiency, it also reduces incident response times, limits operational disruption, and supports consistent decision-making during security events. These benefits contribute directly to service continuity and contractual obligations to clients, aligning cybersecurity activities with broader business and regulatory expectations under DORA and NIS2.

Furthermore, the governance focus of NIST CSF 2.0 encourages regular review of risk appetite and control effectiveness at senior management level. This ensures that security strategy remains adaptive to organisational growth, emerging threats, and evolving regulatory requirements, rather than becoming static or reactive.

7. Risk Management: Limits and Trade-Offs

While the proposed security controls significantly reduce the likelihood and impact of cyber incidents, they cannot eliminate risk entirely. This limitation is particularly evident in the area of software supply chain security, which remains a complex and rapidly evolving challenge. ResilientIntellect relies on third-party libraries, open-source components, and external development tools, all of which introduce dependencies that are not fully under the organisation's direct control. Even with measures such as dependency scanning, code signing, and secure build pipelines, the risk of zero-day vulnerabilities or upstream compromises cannot be completely mitigated [8].

Recognising these constraints is an important aspect of effective risk management. Rather than pursuing absolute prevention, which is neither practical nor cost-effective, the proposed programme prioritises risk visibility, early detection, and response capability. This includes maintaining clear visibility of software dependencies, monitoring for anomalous behaviour within development and deployment environments, and ensuring that incident response processes are well-rehearsed and proportionate to the potential impact of an event.

This approach aligns with contemporary risk management principles and regulatory expectations under DORA and NIS2, which emphasise resilience and the ability to respond to and recover from incidents rather than the unrealistic assumption that all threats can be prevented [1][2]. By explicitly acknowledging trade-offs between security, operational efficiency, and innovation, ResilientIntellect can make informed decisions about where to accept, mitigate, or transfer risk in support of its broader business objectives.

8. Disaster Recovery and Ransomware Scenarios

Ransomware attacks targeting source code repositories and AI training datasets present a dual threat to ResilientIntellect: immediate operational disruption and long-term loss of intellectual property. Given the organisation's reliance on proprietary AI models and continuously evolving threat intelligence, the impact of such an incident would extend beyond short-term service outages to affect competitive advantage and client confidence.

The proposed Disaster Recovery Plan therefore prioritises the restoration of development environments alongside client-facing services. This reflects the reality that prolonged unavailability of source code or training data would significantly delay service delivery and innovation. The use of immutable backups, combined with segregated recovery environments, reduces the likelihood of ransomware persistence or reinfection during the restoration process. These measures are particularly important in hybrid cloud environments, where shared infrastructure can otherwise amplify the spread of an attack.

In addition to technical recovery capabilities, regular disaster recovery testing plays a critical governance role. Scheduled recovery exercises and simulated ransomware scenarios provide assurance that recovery objectives are achievable within acceptable timeframes. They also generate documented evidence of operational resilience, which can be presented to regulators and auditors to demonstrate compliance with DORA's resilience and testing requirements [1]. Through this approach, disaster recovery is positioned not merely as a technical safeguard, but as an integral component of organisational governance and regulatory accountability.

9. Protection Mechanisms and Technical Controls

The protection mechanisms selected for ResilientIntellect's EU unit are designed to support a Zero-Trust-aligned operating model while remaining scalable and manageable as the organisation grows. Rather than relying on isolated security tools, the proposed approach emphasises integrated and automated controls that provide visibility across cloud infrastructure, endpoints, identities, and workloads.

Cloud Security Posture Management (CSPM) tools play a central role in securing the organisation's cloud environments by continuously monitoring configurations against approved baselines and regulatory requirements. Given the dynamic nature of cloud services, CSPM enables early identification of misconfigurations that could otherwise expose sensitive AI training data or internal systems. Where appropriate, automated remediation reduces reliance on manual intervention and supports consistent enforcement of security policies across environments [6].

Security Information and Event Management (SIEM), combined with Security Orchestration, Automation, and Response (SOAR), provides centralised visibility and coordinated response capabilities. SIEM aggregates logs and telemetry from cloud platforms, endpoints, and AI development environments, enabling correlation of events that might otherwise go unnoticed. SOAR builds on this capability by automating routine response actions, such as isolating compromised endpoints or revoking access credentials. While automation improves response speed and consistency, human oversight is retained for high-impact incidents, ensuring that automated actions remain proportionate and context-aware.

Endpoint Detection and Response (EDR) is particularly important in a hybrid working environment where endpoints frequently operate outside traditional network boundaries. By continuously monitoring endpoint behaviour, EDR tools enable rapid detection and containment of malicious activity, reducing the likelihood of lateral movement within the organisation's environment. This aligns with Zero Trust principles by treating endpoints as dynamic risk factors rather than trusted assets.

Finally, phishing-resistant multi-factor authentication (MFA) strengthens identity assurance by reducing reliance on knowledge-based credentials, which remain a common attack vector. Hardware-backed or cryptographic MFA mechanisms significantly reduce the risk of credential compromise and support regulatory expectations for strong access controls. Together, these protection mechanisms form a layered defence that balances security effectiveness with operational practicality, supporting both regulatory compliance and day-to-day business operations.

10. Conclusion

This expanded Zero-Trust-aligned security programme demonstrates how ResilientIntellect Ltd can meet its regulatory obligations under DORA and NIS2 while continuing to support innovation, scalability, and sustainable business growth. By embedding governance, risk management, and operational resilience into organisational structures, policy frameworks, and technical controls, the proposed design moves beyond a narrow, compliance-driven view of security towards a more strategic and integrated security posture.

Importantly, the programme recognises that effective cybersecurity is not solely defined by the number of controls implemented, but by how well those controls align with organisational objectives and risk tolerance. By explicitly acknowledging the limitations of current security practices, particularly in areas such as software supply chain security and cloud environments, the design adopts a pragmatic, risk-based approach rather than pursuing unrealistic prevention guarantees. This balance between security control and operational flexibility enables ResilientIntellect to adapt to emerging threats

and regulatory changes, positioning the organisation to operate effectively within an increasingly complex and dynamic threat landscape.

References

[1] European Union, Regulation (EU) 2022/2554 of the European Parliament and of the Council on Digital Operational Resilience for the Financial Sector (DORA), Official Journal of the European Union, Dec. 2022.

[2] European Union, Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity Across the Union (NIS2 Directive), Official Journal of the European Union, Dec. 2022.

[3] International Organization for Standardization, ISO/IEC 27001:2022 – Information Security, Cybersecurity and Privacy Protection – Information Security Management Systems, ISO, Geneva, Switzerland, 2022.

[4] National Institute of Standards and Technology, The NIST Cybersecurity Framework (CSF) 2.0, NIST, Gaithersburg, MD, USA, 2024.

[5] D. Ribeiro, M. Santos, and L. Chen, “Toward Effective AI Governance: A Review of Principles,” arXiv preprint arXiv:2505.23417, May 2025.

[6] European Commission, Ethics Guidelines for Trustworthy AI, High-Level Expert Group on Artificial Intelligence, Brussels, Belgium, 2019.

[7] J. Kindervag, “Build Security Into Your Network’s DNA: The Zero Trust Network Architecture,” Forrester Research, Cambridge, MA, USA, 2010.

[8] S. Boyson, “Cyber Supply Chain Risk Management: Revolutionizing the Strategic Control of Critical IT Systems,” Technovation, vol. 34, no. 7, pp. 342–353, 2014.

[9] International Organization for Standardization, ISO/IEC 42001:2023 – Artificial Intelligence Management Systems, ISO, Geneva, Switzerland, 2023.

[10] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST, Gaithersburg, MD, USA, 2023.

[11] Gartner, Market Guide for Zero Trust Network Access, Gartner Research, Stamford, CT, USA, 2023.

[12] OECD, AI Governance and the Role of Human Oversight, Organisation for Economic Co-operation and Development, Paris, France, 2021.

[13] M. Alshammari and A. Simpson, “Security Orchestration, Automation and Response (SOAR): A Systematic Review,” *Computers & Security*, vol. 105, 2021.

assignment 1 tyler

GRADEMARK REPORT

FINAL GRADE

GENERAL COMMENTS

57 /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)

TYLER BATE

57

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)



You have engaged with current, peer-reviewed academic literature, standards, and regulatory sources (ISO, NIST, EU AI Act). Literature is synthesised effectively and used to support critical argumentation.

17

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 very good understanding of AI governance, ethics, lifecycle risk, and regulatory compliance. Concepts are articulated clearly and accurately.

8

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)



Critical analysis across AI governance, ethics, and resilience. Strong evaluation of trade-offs, regulatory tensions, and implementation challenges. You have discussion task 2 appropriately.

20

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)



You have not provided any diagrams, tables, charts and any other design methods utilise will enhance or provide marks in this section. Task 2 could at least provide an organisational chart.

0

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)



Report is readable but you have not use the IEEE two column report format for at least task 1. Paragraphs are not fully justified. Overall it could be more presentable.

12