

Data Security and Compliance ACCA5034

SECURING THE MODERN ENTERPRISE: FROM ZERO TRUST TO
AI GOVERNANCE

Daniel Mark Jones
0300760

Contents

Task 1: Implementing Zero Trust Architecture in the Modern Enterprise: Principles, Challenges and Technological Enablers (2026)..... 3

Task 2 – Proposal and Implementation Plan for the Information Security Unit at Aegis FinTech Innovations Ltd. 7

 Executive Summary..... 7

 INTRODUCTION AND STRATEGIC CONTEXT 7

 WORK-STREAM 1: ORGANISATIONAL DESIGN AND RESOURCING..... 8

 Proposed Organisational Structure 8

 WORK-STREAM 2: POLICY FRAMEWORK DEVELOPMENT 9

 Acceptable Use Policy for AI and Cloud Services..... 9

 Hybrid Work and Remote Access Security Policy 9

 Third-Party and Cloud Vendor Risk Management Policy..... 9

 WORK-STREAM 3: REGULATORY COMPLIANCE STRATEGY 9

 WORK-STREAM 4: TECHNICAL STANDARD AND ARCHITECTURE 10

 Core Security Framework 10

 Proposed Zero Trust Architecture 11

 WORK-STREAM 5: STRATEGIC RISK ASSESSMENT..... 11

 WORK-STREAM 6: INCIDENT RESPONSE AND BUSINESS CONTINUITY 12

 WORK-STREAM 7: GOVERNANCE AND STRATEGIC ROADMAP 13

 Cyber Security Steering Committee..... 13

CONCLUSION..... 13

References..... 14

Task 1: Implementing Zero Trust Architecture in the Modern Enterprise: Principles, Challenges and Technological Enablers (2026)

Daniel Jones 0300760
Computing Portfolio – ACCA5034
University of Wales Trinity St David

Abstract – The growth of cloud computing, hybrid working, mobile access and third-party digital services has reduced the effectiveness of security models based on a trusted internal network perimeter. This report critically evaluates Zero Trust Architecture as an approach to strengthening enterprise resilience against identity compromise, lateral movement, and data breaches in hybrid-cloud environments. The research draws upon recent peer-reviewed literature, government guidance, industry threat reports and official investigations of the Change Healthcare ransomware attack and Microsoft Exchange Online intrusion. The findings indicate that Zero Trust can improve security through explicit verification, least-privilege access, identity governance, microsegmentation, cloud-native policy enforcement and continuous monitoring. However, its effectiveness is constrained by legacy systems, incomplete asset visibility, organisational cost, privacy concerns and the potential compromise of identity and policy infrastructure. These findings inform a practical proposal for establishing an Information Security Unit at Aegis FinTech Innovations Ltd. The proposal combines the NIST Cybersecurity Framework 2.0 with a phased Zero Trust Architecture. It prioritises identity security, cloud governance, GDPR and DORA readiness, third-party risk management, ransomware-resistant backups, and incident-response capability. Overall, the report concludes that Zero Trust can substantially improve organisational resilience when implemented as a phased, risk-led programme integrating governance, people, processes, and technology rather than as a standalone security product.

I. INTRODUCTION

Modern organisations no longer operate within clearly defined network boundaries. Employees connect from home and mobile devices, business applications are distributed across public clouds and software-as-a-service platforms, and suppliers frequently require direct access to organisational data and systems. Under these conditions, physical location or connection to an internal network provides limited evidence that a user, device, or application should be trusted. The UK National Cyber Security Centre describes Zero Trust as an architectural approach in which inherent network trust is removed, the network is treated as hostile, and each access request is verified against policy [1]. NIST similarly defines Zero Trust as a shift from static network-perimeter defence towards the protection of users, assets, and individual resources [2]. Academic literature nevertheless indicates that there is no completely uniform understanding of Zero Trust, with definitions frequently shaped by vendors, government agencies, and organisational requirements [3].

This paper addresses the following research question:

To what extent can Zero Trust Architecture improve enterprise resilience against identity compromise,

lateral movement and data breaches in hybrid-cloud environments, and what barriers limit its successful implementation?

It argues that Zero Trust can significantly limit unauthorised access and the potential impact of a compromise. However, Zero Trust should be treated as a long-term security and governance programme rather than a product or guarantee against cyberattack.

II. RESEARCH APPROACH

This paper uses a critical literature review combining peer-reviewed research published from 2020 onwards with government guidance, current industry threat reports, and official case-study evidence. NIST and NCSC publications are used to define Zero Trust principles, while academic studies are used to compare implementation models and identify limitations.

Industry reports provide recent evidence about attack methods but are treated cautiously because security vendors may have commercial interests in promoting particular technologies. Official investigations and government hearings are used for the case studies because they provide stronger evidence than news articles or vendor blogs.

III. ZERO TRUST PRINCIPLES AND ARCHITECTURE

Zero Trust reverses the assumption of internal trust. Under NIST SP 800-207, access to a resource should be granted on a per-session basis, determined through dynamic policy, and limited to the minimum privileges required [2]. Authentication and authorisation occur before access is established, while behavioural, environmental and device information may be used to reassess access during a session. Zero Trust therefore does not mean that trust is permanently absent; instead, confidence must be established for a specific request and continually reconsidered.

The logical NIST architecture contains three particularly important components. The Policy Engine decides whether access should be granted. The Policy Administrator establishes or terminates the communication path, while the Policy Enforcement Point enables, monitors, and terminates the connection between a subject and a resource [2]. Inputs may include user identity, device health, threat intelligence, location, previous behaviour, and the sensitivity of the requested data.

The NCSC and NIST approaches are complementary. NIST provides a detailed logical architecture, whereas the NCSC offers principles focused on practical enterprise design. The NCSC emphasises understanding users, devices, services,

and data; assessing device and user behaviour; using policy to authorise requests; and authenticating and authorising throughout the architecture [1]. Both frameworks reject trust based solely on network location.

However, Zero Trust should not be reduced to the slogan “never trust, always verify.” Buck et al. found that Zero Trust encompasses technological, organisational, and strategic considerations, while also identifying gaps between academic research and practitioner claims [3]. An organisation can deploy multifactor authentication or network segmentation without having a complete Zero Trust Architecture. Effective implementation requires coordinated controls across identities, endpoints, networks, applications, and data.

IV. MODERN ENTERPRISE THREAT LANDSCAPE

The continued importance of compromised credentials strengthens the business case for Zero Trust, as well as exploited vulnerabilities and ransomware. Verizon's 2026 Data Breach Investigations Report states that software vulnerabilities became the leading initial access route in its dataset, ahead of stolen passwords [4]. IBM's 2026 X-Force reporting similarly identifies increased exploitation of public-facing applications, expanding supply-chain risk and persistent weaknesses in foundational controls [5].

These findings demonstrate that Zero Trust must address more than human identities. An attacker may gain access by compromising a user account, exploiting an internet-facing application, stealing an API token, or abusing a trusted supplier. Workloads, applications, and automated services therefore require identities and access policies in the same way as human users.

Identity compromise remains particularly important because valid credentials may allow malicious activity to resemble legitimate access. Password reuse, phishing, token theft, and weakly protected privileged accounts can allow an attacker to bypass perimeter controls. Once inside a flat network, the attacker may conduct discovery, escalate privileges, and move laterally towards valuable data.

Zero Trust cannot prevent every vulnerability from being exploited. Its contribution is to reduce the permissions available after the initial compromise. Application-specific access, device validation, and segmentation can prevent one compromised identity or workload from automatically reaching unrelated resources. Continuous monitoring may also identify changes in behaviour after legitimate authentication has occurred.

V. TECHNOLOGICAL ENABLERS

A. Identity Governance and Authentication

Identity is a central Zero Trust control point. Organisations require reliable processes for creating, changing, and removing accounts, including joiner, mover, and leaver procedures. Access should be based on business need, reviewed regularly, and removed when it is no longer required. Privileged accounts should be separated from normal user accounts and controlled through Privileged Access Management.

Multifactor authentication increases resistance to password theft, although its strength depends upon the method used. Phishing-resistant authenticators and device-bound credentials provide stronger protection than reusable passwords and

one-time codes. Authentication should also be supported by conditional-access policies that consider device compliance, user risk, location, and the requested resource.

B. ZTNA, SASE and Microsegmentation

Zero Trust Network Access provides access to specifically authorised applications instead of exposing a wider internal network. This makes ZTNA particularly useful for remote employees and third parties. It can reduce the ability of a compromised endpoint to discover or directly contact unrelated internal services.

However, ZTNA should not be presented as equivalent to Zero Trust. It is one enforcement mechanism within a broader architecture. Similarly, Secure Access Service Edge can combine ZTNA with cloud-delivered security capabilities, but purchasing a SASE platform does not resolve weak identity governance, poor data classification, or excessive permissions.

Microsegmentation divides an environment into smaller security zones or workload groups. Policies can then limit communication between production, development, database, and administrative systems. This reduces the potential blast radius of a compromise because the attacker cannot assume that control of one endpoint provides access to all other systems.

C. Cloud-Native Enforcement

Hybrid and multi-cloud systems introduce additional complexity because applications may consist of containers, APIs, serverless functions and microservices distributed across several providers. NIST SP 800-207A extends Zero Trust principles to cloud-native applications and discusses granular policy enforcement through components such as API gateways, service meshes and application identities [6].

NIST's practical implementation guide demonstrates that Zero Trust can involve identity governance, microsegmentation, SASE and software-defined perimeter approaches [7]. It also shows that different technical designs can implement similar principles. This supports the view that Zero Trust is an architectural outcome rather than a single prescribed product.

D. Monitoring, Analytics and Automated Response

Continuous verification depends on visibility. Identity logs, endpoint telemetry, cloud configuration information, application events, and network activity must be collected and analysed. SIEM platforms can correlate these signals, while EDR and XDR tools provide endpoint and workload detection.

Yeoh et al. identify analytics, visibility, automation, and orchestration as important dimensions of Zero Trust maturity [8]. Automated responses might include requiring additional authentication, terminating a session, disabling a credential, or isolating an endpoint. Such automation can reduce response time, especially when an attacker moves faster than a human security team can investigate.

Automation also introduces risk. Cao et al. note that automated Zero Trust systems must process large quantities of heterogeneous data and make policy decisions without creating unacceptable errors or operational disruption [13]. A false positive could interrupt a critical business process, while an inaccurate risk model might allow malicious access. Human oversight and tested fallback procedures therefore remain necessary.

VI. IMPLEMENTATION CHALLENGES AND LIMITATIONS

A. Legacy Systems and Asset Visibility

Zero Trust relies on knowledge of organisational identities, devices, applications, data, and communication paths. Many organisations lack a complete asset inventory, particularly where departments have independently adopted cloud services. Unknown resources cannot be protected through accurate access policies.

Legacy systems may not support modern authentication, encryption, granular authorisation, or security telemetry. Replacing these systems may be expensive or operationally disruptive. Organisations may consequently require temporary gateways or compensating controls, but these can increase architectural complexity.

Phiayura and Teerakanok describe migration as an enterprise-wide process rather than an immediate replacement of existing security [9]. Itodo and Ozer similarly identify implementation challenges involving technology, governance, people, and organisational processes [10]. A phased migration is therefore more realistic than an attempt to deploy every Zero Trust capability simultaneously.

B. Cost, Skills, and Organisational Change

Security controls may face resistance if users experience frequent authentication prompts or lose broad access previously considered convenient. Poorly designed controls may encourage employees to seek unofficial workarounds. Implementation should therefore balance security with usability and involve application owners, employees, legal teams, and business managers. Economic analysis also indicates that Zero Trust investment can reduce expected risk impact, although the result depends on organisational scale and implementation assumptions [23].

C. Privacy and Control-Plane Risk

Continuous assessment can involve collecting information about user location, device condition, and behaviour. This creates privacy and employee-monitoring concerns, particularly where data is retained without a clear purpose. A recent multidimensional survey similarly identifies unresolved security and privacy questions around continuous identity, device, and contextual assessment [24]. Organisations must therefore define proportionate collection, access, and retention rules rather than assuming that more monitoring always produces better security.

Fernandez and Brazhuk caution against treating Zero Trust as a completely new or universally sufficient architecture, noting that many of its controls originate from established security concepts such as least privilege, defence in depth and strong authentication [11]. This criticism is useful because

vendor marketing can imply that Zero Trust eliminates the need for other security disciplines.

Ferretti et al. identify a further weakness: Zero Trust architectures frequently assume that the control plane itself remains trustworthy and available [12]. An identity provider, policy engine or central certificate authority can become a high-value target and potential single point of failure. Zero Trust components therefore require redundancy, privileged-access protection, secure key management, and independent monitoring.

VII. REAL-WORLD CASE STUDIES

A. Change Healthcare Ransomware Attack

The 2024 Change Healthcare attack demonstrates the consequences of weak remote-access security. Attackers used compromised credentials to access a remote-access service that did not have multifactor authentication enabled. After entering the environment, they moved laterally, removed data and later deployed ransomware [14].

MFA could have made the stolen credentials harder to use, but it would not have provided complete protection. Zero Trust would also assess device management, behavioural risk, and the account's business need. Application-specific access, segmentation, and better telemetry could have restricted lateral movement and accelerated detection.

B. Microsoft Exchange Online Intrusion

The Cyber Safety Review Board investigated the 2023 Microsoft Exchange Online intrusion conducted by the actor known as Storm-0558. The investigation concluded that the intrusion was enabled by a sequence of security failures and identified weaknesses involving identity infrastructure, signing keys, token validation, and security logging [15].

This case exposes an important limitation of identity-centred security: the systems that issue and validate identities must themselves be secure. If a trusted signing key or identity provider is compromised, an attacker may create tokens that appear legitimate. Continuous verification cannot be reliable when the evidence used by the policy engine has been corrupted.

The incident therefore supports the survivability concerns identified by Ferretti et al. [12]. Identity providers and policy systems should be isolated, monitored and designed for secure recovery. Cryptographic keys require strong lifecycle management, and customers need sufficient logging to investigate abnormal cloud activity.

VIII. CRITICAL DISCUSSION AND RECOMMENDED APPROACH

The evidence indicates that Zero Trust can improve enterprise resilience, but its benefits are conditional. It reduces dependence on network location; limits access to individual resources and supports faster detection of abnormal behaviour. These capabilities can restrict lateral movement following credential theft or vulnerability exploitation.

Zero Trust cannot remove all cyber risk or guarantee that software, employees, and suppliers remain uncompromised. It may also concentrate risk in identity and policy infrastructure. Its defensible benefit is improved containment and management of compromise, not the elimination of compromise.

A phased implementation should begin with governance and visibility. The organisation should identify critical services, sensitive data, user and workload identities, devices, and access paths. It should then strengthen identity lifecycle processes, privileged access, and multifactor authentication.

The next phase should introduce application-specific ZTNA, device-compliance checks and segmentation around high-value systems. Cloud environments should use workload identities, short-lived credentials, and automated configuration assessment. Identity, endpoint, cloud, and application logs should subsequently be integrated into central monitoring.

IX. CONCLUSION

Zero Trust Architecture provides a more appropriate security model for hybrid enterprises than approaches based on a trusted internal network. By combining explicit verification, least privilege, device assessment, application-specific access, segmentation, and continuous monitoring, it can reduce unauthorised access and limit lateral movement.

However, successful implementation depends on more than deploying ZTNA or multifactor authentication. Organisations require asset visibility, identity governance, resilient policy infrastructure, skilled personnel and sustained executive support. Legacy systems, operational cost, privacy concerns, and control-plane compromise remain significant limitations.

The research question can therefore be answered conditionally: Zero Trust can substantially improve enterprise resilience, but only when implemented as a phased, risk-led programme integrating people, processes, and technology. It should be regarded as a method of reducing and containing cyber risk, not as a promise that all systems and identities will remain secure.

Task 2 – Proposal and Implementation Plan for the Information Security Unit at Aegis FinTech Innovations Ltd.

Executive Summary

Aegis FinTech Innovations Ltd. is preparing to expand its AI-powered financial analytics platform into the European Union. Its reliance on cloud infrastructure, customer financial information, artificial intelligence, APIs, and third-party technology providers creates significant exposure to credential compromise, ransomware, cloud misconfiguration, supply-chain attacks, and regulatory non-compliance.

This proposal recommends establishing a dedicated Information Security Unit governed through the NIST Cybersecurity Framework 2.0 and supported by a phased Zero Trust Architecture. NIST CSF 2.0 provides a risk-management structure based on the Govern, Identify, Protect, Detect, Respond and Recover functions. In contrast, Zero Trust provides the technical principles for identity, device, application, and data access [2], [16].

The proposed 12-month programme prioritises security governance, asset visibility, identity management, cloud security, GDPR and DORA readiness, third-party risk, ransomware-resistant backups, and incident-response capability.

INTRODUCTION AND STRATEGIC CONTEXT

Aegis FinTech's planned European expansion will increase the volume and sensitivity of the personal, financial, and analytical data it processes. Its hybrid-cloud operating model also means that users, devices, applications, and data are no longer contained within a single trusted network perimeter.

Task 1 concluded that Zero Trust is more suitable for hybrid enterprises than perimeter-centred security because access should be based on verified identity, device condition, business need, and resource sensitivity rather than network location [1], [2]. The proposed Information Security Unit will therefore implement least privilege, explicit verification, segmentation, continuous monitoring, and assume-breach principles.

WORK-STREAM 1: ORGANISATIONAL DESIGN AND RESOURCING

Proposed Organisational Structure




Fig. 1. Proposed Organisational Structure

The CISO will own the security strategy, budget, risk reporting, and executive communication. The Cloud Security Architect will establish secure cloud landing zones, segmentation, encryption, and cloud security baselines. The IAM/Zero Trust Engineer will manage identity lifecycle processes, multifactor authentication, privileged access, conditional access, and ZTNA.

The DevSecOps Engineer will integrate code scanning, dependency analysis, secrets detection, infrastructure-as-code assessment, and application testing into development pipelines. This role is particularly important because recent research on cloud-native financial systems identifies Zero Trust controls across the software-development lifecycle as necessary for managing microservice and supply-chain exposure [25]. The GRC and Vendor Risk Analyst will manage policies, risks, audits, supplier assurance, and compliance evidence.

The SOC and Incident Response Lead will manage monitoring, incident investigation, and ransomware response. A managed security service provider should initially provide 24-hour monitoring because establishing a fully staffed internal SOC would be disproportionate for a Series B startup.

During the first year, Aegis should prioritise recruitment of the CISO, Cloud Security Architect, IAM/Zero Trust Engineer, and GRC/Vendor Risk Analyst. SOC monitoring should initially be delivered through an MSSP, while specialist DPO, penetration-testing and incident-response services may be retained externally. The remaining engineering roles can be recruited in phases as the platform, regulatory exposure, and security workload increase.

Aegis should appoint an internal or external Data Protection Officer, subject to confirming whether the appointment is mandatory under GDPR Article 37. The DPO must retain independence and direct access to the highest level of management under Article 38 [17]. The DPO should therefore work closely with the CISO but should not be managed in a manner that creates a conflict of interest.

Task 1 found that Zero Trust requires coordinated capabilities across identity, endpoints, applications, networks, data, analytics, and automation [8]. The proposed structure assigns responsibility for each capability rather than treating Zero Trust as a technology-only project.

WORK-STREAM 2: POLICY FRAMEWORK DEVELOPMENT

Acceptable Use Policy for AI and Cloud Services

The policy will apply to employees, contractors and third parties using cloud platforms, AI tools, models, or datasets on behalf of Aegis. Only approved cloud and AI services may be used. Personal, customer, financial, source-code or commercially sensitive data must not be submitted to public generative AI services unless the service has undergone security, privacy, and contractual approval. Data must be classified before use, and significant AI-generated outputs must receive human review.

Corporate platforms must use centralised identity management, multifactor authentication, and least-privilege access. Model training data, prompts, outputs, and administrative actions must be logged where proportionate. Employees must report unauthorised AI tools, accidental disclosures, and suspicious model behaviour.

The research established that trust should be continuously evaluated using identity, device, resource, and contextual information [1], [2]. The policy consequently prohibits unmanaged identities and anonymous access to business cloud or AI systems.

Hybrid Work and Remote Access Security Policy

Remote access will be restricted to approved users and managed devices. Multifactor authentication, endpoint detection, device encryption, and current security patches will be mandatory.

ZTNA will replace broad VPN access where technically feasible. Users will receive access to specific approved applications rather than general access to the corporate network. Privileged administration will require separate administrator accounts, privileged access management and dedicated secure workstations.

Access may be denied or restricted where the device is unmanaged, security software is disabled, the connection is considered high risk, or the user's behaviour is inconsistent with normal activity.

The Change Healthcare case examined in the research demonstrated the consequences of an externally accessible service without multifactor authentication: compromised credentials enabled initial access and subsequent lateral movement [14].

Third-Party and Cloud Vendor Risk Management Policy

Suppliers will be classified according to the sensitivity of information processed, service criticality and degree of technical access. Critical vendors will undergo due diligence before contract approval and annual reassessment thereafter.

Contracts must address security standards, access controls, encryption, incident notification, audit evidence, subcontractors, service continuity, data location, data deletion, and termination assistance. Critical integrations must use separate machine identities, short-lived credentials, and least-privilege permissions.

Where DORA applies, Aegis will maintain a register of ICT contractual arrangements and assess concentration, subcontracting and exit risks [18].

The research also established that a supplier should not inherit broad trust merely because it has passed an initial assessment. Supplier identities, devices, APIs, and workloads must remain restricted, reassessed, and monitored.

WORK-STREAM 3: REGULATORY COMPLIANCE STRATEGY

Aegis must first obtain legal advice to determine whether it will fall directly within DORA's definition of a financial entity, operate as an ICT third-party provider to regulated financial entities, or become subject to DORA-related contractual requirements from customers.

DORA has applied since 17 January 2025 and covers ICT risk management, incident handling, resilience testing, and ICT third-party risk for regulated financial entities [18]. Table I translates these requirements and the parallel GDPR work programme into a 12-month compliance roadmap.

TABLE I. GDPR AND DORA COMPLIANCE ROADMAP

Period	GDPR activities	DORA-aligned activities	Evidence
Months 1-2	Data inventory, processing records, lawful-basis review and DPO assessment	Applicability assessment, ICT asset inventory and gap analysis	Compliance gap report
Months 3-4	Privacy notices, retention rules, rights procedures and DPIAs	ICT risk framework, critical-service mapping and risk tolerance	Approved frameworks
Months 5-6	Processor contracts, transfer controls and security measures	Incident classification, supplier register and contract review	Vendor and incident registers
Months 7-9	AI-platform DPIA, privacy training and breach exercise	Resilience, backup-restoration and supplier continuity tests	Test reports
Months 10-12	Privacy audit and corrective action	Board assurance review and regulatory evidence pack	Compliance dashboard

The GDPR work programme will prioritise processing records, data minimisation, retention, transparency, data-subject rights, processor contracts, security of processing and breach management. High-risk AI profiling or evaluation should be assessed through a Data Protection Impact Assessment under Article 35 [17].

DORA Articles 5-16 concern ICT governance and risk management; Articles 17-23 address incident management; Articles 24-27 address resilience testing; and Articles 28-44 address ICT third-party risk. Article 11 requires documented backup, restoration and recovery arrangements and periodic testing [18].

WORK-STREAM 4: TECHNICAL STANDARD AND ARCHITECTURE

Core Security Framework

Aegis should adopt NIST Cybersecurity Framework 2.0 as its principal risk-management framework [16]. It is appropriate because it supports executive and technical communication, covers governance through recovery, enables current and target profiles, is technology-neutral, and can be mapped to GDPR, DORA and ISO/IEC 27001 controls.

ISO/IEC 27001 certification may be considered after the first 12 months, but initially seeking certification could divert resources from addressing immediate identity, cloud, and ransomware risks.

Proposed Zero Trust Architecture




Fig. 2. Proposed Zero Trust Architecture for Aegis FinTech.

The Identity Provider will act as the authoritative access-control source. Identity Governance and Administration will automate joiner, mover, and leaver processes, while PAM will restrict privileged credentials.

The policy engine will evaluate identity, device posture, user behaviour, threat intelligence, and resource sensitivity. ZTNA, API gateways and service-mesh controls will act as Policy Enforcement Points [2], [6].

Microsegmentation will separate production, development, AI training, model deployment, databases, and administrative services. CSPM and Cloud Infrastructure Entitlement Management will identify configuration weaknesses and excessive cloud permissions. SIEM and SOAR will correlate identity, endpoint, application and cloud telemetry and support controlled automated responses.

NIST SP 1800-35 demonstrates that identity governance, microsegmentation, SASE and software-defined perimeter approaches can implement Zero Trust in different ways [7].

The research concluded that Zero Trust is an architectural approach rather than a single product. The proposed design therefore combines identity governance, endpoint security, ZTNA, segmentation, cloud-posture management, monitoring, and recovery controls. It also protects the identity and policy infrastructure because compromise of the control plane could undermine all access decisions [12], [15].

WORK-STREAM 5: STRATEGIC RISK ASSESSMENT

TABLE II. PRIORITY RISK REGISTER

Risk	Inherent risk	Principal mitigation	Residual risk	Owner
Credential compromise followed by ransomware and lateral movement	Critical	Phishing-resistant MFA; PAM; ZTNA; EDR; segmentation; behavioural monitoring	Medium	CISO
Cloud misconfiguration or excessive permissions exposing data and AI assets	Critical	CSPM/CIEM; secure cloud baselines; IaC scanning; encryption	Medium	Cloud Architect

Risk	Inherent risk	Principal mitigation	Residual risk	Owner
Third-party or software-supply-chain compromise	High	Due diligence; contractual controls; least-privilege integrations; monitoring; exit plan	Medium	GRC Lead

Table II distinguishes inherent risk from the expected residual position after treatment. Identity compromise is prioritised because the research demonstrated that valid credentials may allow attackers to appear legitimate. Cloud misconfiguration is similarly significant because Aegis's platform depends on cloud data, APIs, and AI workloads. Third-party compromise remains a high inherent risk because suppliers may have extensive technical access or provide essential services.

A secondary risk register should also cover AI data poisoning, model theft, insecure APIs, insider misuse, and failure to restore systems following ransomware.

WORK-STREAM 6: INCIDENT RESPONSE AND BUSINESS CONTINUITY

Aegis will adopt a ransomware response playbook aligned with the Govern, Identify, Protect, Detect, Respond and Recover outcomes in the NIST ransomware community profile [22].

Detection and Activation

The SOC will validate alerts, identify affected identities and systems, determine whether data has been removed and assign an incident severity. The CISO will activate the incident team for confirmed or potentially material incidents.

Containment

Affected endpoints and cloud workloads will be isolated. Compromised accounts, tokens, API keys, and sessions will be disabled or revoked. Network communication between segments will be restricted, while logs and forensic evidence will be preserved.

Eradication

The technical team will remove malware and persistence, patch exploited vulnerabilities, rotate credentials and rebuild compromised systems from approved images. Identity and policy systems will receive particular attention because attackers may attempt to retain privileged access.

Recovery

Critical services will be restored according to agreed recovery-time and recovery-point objectives. Aegis will maintain encrypted immutable backup copies using object-lock or write-once-read-many controls, supported by offline or logically isolated recovery copies and separate backup-administration identities. These measures prevent compromise of production credentials from being used to alter or delete recovery data. The NCSC recommends that backup systems be isolatable and resilient to destructive actions, with restoration from earlier versions remaining possible [21].

Recovery testing will verify data integrity, application functionality and security before systems are reconnected. Increased monitoring will continue until the incident commander approves the return to normal operation.

Communication

The CISO will act as incident commander. The DPO and Legal Counsel will assess privacy and regulatory notifications, while the CEO or nominated communications lead will manage Board, customer, and media communications.

A notifiable personal-data breach must be reported to the appropriate supervisory authority without undue delay and, where feasible, within 72 hours of awareness under GDPR Article 33 [17]. The ICO provides the same 72-hour expectation for qualifying UK personal-data breaches [19].

Post-Incident Review

A root-cause analysis will identify control failures, update the risk register, and assign corrective actions. Results will be reported to the Steering Committee and incorporated into future exercises.

The playbook applies the assume-breach principle established in the research: prevention remains necessary, but Aegis must also restrict lateral movement, detect compromise, and recover safely.

WORK-STREAM 7: GOVERNANCE AND STRATEGIC ROADMAP

Cyber Security Steering Committee

The committee should include the CEO, CISO, CTO, CFO, DPO or Legal Counsel, Head of Engineering, Head of Product/AI, Head of Operations and an HR representative. It will meet monthly during the first year and will approve security strategy, risk appetite, policies, major exceptions, and investment priorities. It will also review incidents, compliance progress, KPIs and overdue remediation. Table III converts these governance priorities into quarterly deliverables and measurable success criteria.

TABLE III. 12-MONTH STRATEGIC ROADMAP

Quarter	Priority	Key deliverables	Success measures
Q1	Govern and discover	Appoint roles, create committee, complete asset and data inventories, conduct risk and compliance gaps	100% of critical services identified, classified, and assigned owners
Q2	Identity and policy	Deploy MFA and PAM, approve three policies, establish vendor and processing registers	100% of privileged accounts protected by phishing-resistant MFA and PAM
Q3	Protect and detect	Implement ZTNA pilot, EDR, CSPM, SIEM use cases and initial segmentation	At least 95% managed-endpoint compliance
Q4	Respond and assure	Ransomware exercise, restore test, penetration test, GDPR/DORA review, and Board assurance	100% of critical services restored within approved RTO and RPO targets

Additional targets will include 100% of critical suppliers assessed before production access; 95% of critical vulnerabilities remediated within the approved service-level target; 100% of critical services included in quarterly restoration tests; mean time to contain high-severity incidents below four hours; 100% of access reviews completed on time; and 100% annual security-training completion. The Steering Committee will review these measures monthly.

FCA operational-resilience guidance emphasises identifying important business services, setting impact tolerances and testing vulnerabilities. These principles should be adopted even before Aegis's final regulatory status is confirmed [20].

The research showed that organisations frequently encounter difficulties when they begin with technology before establishing visibility, ownership, and governance [9], [10]. The roadmap therefore starts with governance and asset discovery before introducing advanced enforcement and automation.

CONCLUSION

Aegis FinTech requires a security function capable of supporting rapid growth, EU expansion, and investor assurance without creating disproportionate cost. NIST CSF 2.0 provides the governance and risk framework, while Zero Trust provides the architectural basis for identity, device, workload, and data protection.

The proposed Information Security Unit combines internal leadership and engineering expertise with initially outsourced monitoring. Its first-year roadmap prioritises identity security, cloud visibility, regulatory readiness, ransomware resilience, and supplier assurance. These recommendations translate the findings of Task 1 into an actionable programme and establish a defensible foundation for GDPR and DORA compliance.

References

- [1] National Cyber Security Centre, "Zero trust architecture design principles," Jan. 16, 2026. [Online]. Available: <https://www.ncsc.gov.uk/col-lection/zero-trust/architecture-design-principles>. Accessed: Jul. 20, 2026.
- [2] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," National Institute of Standards and Technology, NIST Special Publication 800-207, Aug. 2020, doi: 10.6028/NIST.SP.800-207. Available: <https://csrc.nist.gov/pubs/sp/800/207/final>
- [3] C. Buck, C. Olenberger, A. Schweizer, F. Völter, and T. Eymann, "Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero-trust," *Computers & Security*, vol. 110, Art. no. 102436, Nov. 2021, doi: 10.1016/j.cose.2021.102436. Available: <https://doi.org/10.1016/j.cose.2021.102436>
- [4] Verizon, "2026 Data Breach Investigations Report," 2026. [Online]. Available: <https://www.verizon.com/business/resources/reports/dbir/>. Accessed: Jul. 20, 2026.
- [5] IBM Security, "X-Force Threat Intelligence Index 2026," 2026. [Online]. Available: <https://www.ibm.com/reports/threat-intelligence>. Accessed: Jul. 20, 2026.
- [6] R. Chandramouli and Z. Butcher, "A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Cloud Environments," NIST Special Publication 800-207A, Sept. 2023, doi: 10.6028/NIST.SP.800-207A. Available: <https://csrc.nist.gov/pubs/sp/800/207/a/final>
- [7] National Cybersecurity Center of Excellence, "Implementing a Zero Trust Architecture," NIST Special Publication 1800-35, June 2025. Available: <https://csrc.nist.gov/pubs/sp/1800/35/final>
- [8] W. Yeoh, M. Liu, M. Shore, and F. Jiang, "Zero trust cybersecurity: Critical success factors and a maturity assessment framework," *Computers & Security*, vol. 133, Art. no. 103412, Oct. 2023, doi: 10.1016/j.cose.2023.103412. Available: <https://doi.org/10.1016/j.cose.2023.103412>
- [9] P. Phiyayura and S. Teerakanok, "A comprehensive framework for migrating to Zero Trust Architecture," *IEEE Access*, vol. 11, pp. 19487-19511, 2023, doi: 10.1109/ACCESS.2023.3248622. Available: <https://doi.org/10.1109/ACCESS.2023.3248622>
- [10] C. Itodo and M. Ozer, "Multivocal literature review on zero-trust security implementation," *Computers & Security*, vol. 141, Art. no. 103827, June 2024, doi: 10.1016/j.cose.2024.103827. Available: <https://doi.org/10.1016/j.cose.2024.103827>
- [11] E. B. Fernandez and A. Brazhuk, "A critical analysis of Zero Trust Architecture (ZTA)," *Computer Standards & Interfaces*, vol. 89, Art. no. 103832, Apr. 2024, doi: 10.1016/j.csi.2024.103832. Available: <https://doi.org/10.1016/j.csi.2024.103832>
- [12] L. Ferretti, F. Magnanini, M. Andreolini, and M. Colajanni, "Survivable zero trust for cloud computing environments," *Computers & Security*, vol. 110, Art. no. 102419, Nov. 2021, doi: 10.1016/j.cose.2021.102419. Available: <https://doi.org/10.1016/j.cose.2021.102419>
- [13] Y. Cao, S. R. Pokhrel, Y. Zhu, R. Doss, and G. Li, "Automation and orchestration of Zero Trust Architecture: Potential solutions and challenges," *Machine Intelligence Research*, vol. 21, pp. 294-317, 2024, doi: 10.1007/s11633-023-1456-2. Available: <https://doi.org/10.1007/s11633-023-1456-2>
- [14] U.S. Senate Committee on Finance, "Hacking America's Health Care: Assessing the Change Healthcare Cyber Attack and What's Next," hearing record and witness testimony, May 1, 2024. [Online]. Available: <https://www.finance.senate.gov/hearings/hacking-americas-health-care-assessing-the-change-healthcare-cyber-attack-and-whats-next>. Accessed: Jul. 20, 2026.
- [15] Cyber Safety Review Board, "Review of the Summer 2023 Microsoft Exchange Online Intrusion," Mar. 2024. [Online]. Available: <https://www.cisa.gov/resources-tools/resources/CSRB-Review-Summer-2023-MEO-Intrusion>. Accessed: Jul. 20, 2026.
- [16] National Institute of Standards and Technology, "The NIST Cybersecurity Framework (CSF) 2.0," NIST CSWP 29, Feb. 2024, doi: 10.6028/NIST.CSWP.29. Available: <https://csrc.nist.gov/pubs/cswp/29/the-nist-cybersecurity-framework-csf-20/final>
- [17] European Parliament and Council, "Regulation (EU) 2016/679: General Data Protection Regulation," *Official Journal of the European Union*, Apr. 27, 2016. Available: <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>
- [18] European Parliament and Council, "Regulation (EU) 2022/2554 on digital operational resilience for the financial sector," *Official Journal of the European Union*, Dec. 14, 2022. Available: <https://eur-lex.europa.eu/eli/reg/2022/2554/oj/eng>
- [19] Information Commissioner's Office, "Personal data breaches: a guide," Aug. 20, 2025. [Online]. Available: <https://ico.org.uk/for-organisations/report-a-breach/personal-data-breach/personal-data-breaches-a-guide/>. Accessed: Jul. 22, 2026.
- [20] Financial Conduct Authority, "Operational resilience." [Online]. Available: <https://www.fca.org.uk/firms/operational-resilience>. Accessed: Jul. 20, 2026.
- [21] National Cyber Security Centre, "Ransomware-resistant backups," Nov. 22, 2024. [Online]. Available: <https://www.ncsc.gov.uk/collec-tion/ransomware-resistant-backups>. Accessed: Jul. 20, 2026.
- [22] M. Souppaya, K. Scarfone, and D. Dodson, "Ransomware Risk Management: A Cybersecurity Framework 2.0 Community Profile," NIST Interagency/Internal Report 8374 Rev. 1, June 2026, doi: 10.6028/NIST.IR.8374r1. Available: <https://csrc.nist.gov/pubs/ir/8374/r1/final>

- [23] Z. Adahman, A. W. Malik, and Z. Anwar, "An analysis of zero-trust architecture and its cost-effectiveness for organizational security," *Computers & Security*, vol. 122, Art. no. 102911, Nov. 2022, doi: 10.1016/j.cose.2022.102911. Available: <https://doi.org/10.1016/j.cose.2022.102911>
- [24] M. A. Azad, S. Abdullah, J. Arshad, H. Lallie, and Y. H. Ahmed, "Verify and trust: A multidimensional survey of zero-trust security in the age of IoT," *Internet of Things*, vol. 27, Art. no. 101227, Oct. 2024, doi: 10.1016/j.iot.2024.101227. Available: <https://doi.org/10.1016/j.iot.2024.101227>
- [25] D. Shin, J. Kim, I. W. A. J. Pawana, and I. You, "Enhancing cloud-native DevSecOps: A Zero Trust approach for the financial sector," *Computer Standards & Interfaces*, vol. 93, Art. no. 103975, Apr. 2025, doi: 10.1016/j.csi.2025.103975. Available: <https://doi.org/10.1016/j.csi.2025.103975>