

Implementing a Zero Trust Architecture: Principles, Challenges and Technological Enablers

DATA SECURITY AND COMPLIANCE

Research and Case Scenario Report

Implementing a Zero Trust Architecture:

Principles, Challenges and Technological Enablers

IEEE-style in-text citations and reference list

Table of Contents

List of Figures and Tables	3
List of Acronyms	3
Task 1 – Research in Data Security and Compliance	4
1.1 Abstract	4
1.2 Introduction	4
1.3 The Fundamental Principles and Logical Structure	4
1.4 Technological Enablers	5
1.5 Benefits, Challenges and Limitations.....	6
1.6 Real-World Evidence and Case Studies	7
1.7 Recommended Implementation Approach	7
1.8 Conclusion.....	8
Task 2 – CISO Proposal and Implementation Plan	9
2.1 Executive Summary	9
2.2 Work-Stream 1: Organisational Design and Resourcing	9
2.3 Work-Stream 2: Policy Framework Development	10
2.4 Work-Stream 3: Regulatory Compliance Strategy.....	11
2.5 Work-Stream 4: Technical Standards and Architecture	11
2.6 Work-Stream 5: Strategic Risk Assessment	12
2.7 Work-Stream 6: Incident Response and Business Continuity	12
2.8 Work-Stream 7: Governance and the 12-Month Strategic Roadmap	13
2.9 Conclusion.....	12
References	13-14

List of Figures and Tables

Figures	Tables
Figure 1. Context-aware Zero Trust decision cycle	Table 1. Zero Trust technology enablers and limitations
Figure 2. Proposed Information Security Unit	Table 2. Evidence-to-judgement matrix
Figure 3. High-level Zero Trust architecture for Aegis FinTech	Table 3. Security unit roles and initial resourcing
Figure 4. Ransomware response and recovery lifecycle	Table 4. Executive summaries of the three mandatory policies
	Table 5. GDPR and DORA compliance roadmap
	Table 6. High-level strategic risk register
	Table 7. Cyber Security Steering Committee
	Table 8. Twelve-month implementation roadmap and KPIs

List of Acronyms

CISO	Chief Information Security Officer	CSPM	Cloud Security Posture Management
DORA	Digital Operational Resilience Act	DPO	Data Protection Officer
EDR	Endpoint Detection and Response	GDPR	General Data Protection Regulation
IAM / IGA	Identity and Access / Identity Governance and Administration	MFA	Multi-Factor Authentication
PAM	Privileged Access Management	SIEM / XDR	Security Information and Event Management / Extended Detection and Response
SASE	Secure Access Service Edge	ZTA / ZTNA	Zero Trust Architecture / Zero Trust Network Access

TASK 1 - RESEARCH IN DATA SECURITY AND COMPLIANCE

Implementing a Zero Trust Architecture: Principles, Challenges and Technological Enablers

1.1 Abstract

Zero Trust Architecture (ZTA) removes automatic trust based on network location and verifies every access request. This paper examines its principles, challenges and key technologies, including SASE, ZTNA and identity governance. ZTA can reduce excessive access, lateral movement and remote-working risks, but it depends on strong identity controls, asset visibility, device management and monitoring. A phased approach should begin with identity, critical assets and security monitoring before introducing dynamic access and automation.

1.2 Introduction

Modern organisations rely on cloud services, remote devices, partner networks and APIs, making the traditional trusted internal network unreliable. Compromised accounts, unmanaged devices and vulnerable VPNs can allow attackers to enter and move across systems. NIST states that zero trust protects users, assets and resources without trusting them based on location or ownership [1]. Similarly, the NCSC advises treating the network as hostile and verifying every access request [2].

This paper evaluates ZTA principles and technologies, including identity governance, ZTNA, SASE, micro-segmentation and monitoring. It argues that ZTA is most effective as a long-term security program, rather than simply a collection of separate security tools.

1.3 The fundamental principles and logical structure

The basic rule is to make explicit verification. When making access decisions, several factors should be taken into account, such as the user's identity, the condition of the device, the user's location, the resource being requested, the sensitivity of the data, the user's behaviour, and the

up-to-date threat intelligence. Authentication should not be limited to a single login. The policy decision can be examined during the course of a session and access can be denied or revoked if the risk level changes. This approach is more suitable for cloud and hybrid environments than relying on the user's IP address or corporate LAN connection.

The second principle is least privilege, where users get only necessary access for specific tasks and durations, using role-based access control, attribute rules, privileged management, and just-in-time elevation. Least privilege reduces damage from stolen credentials and insider misuse, preventing network exposure from a single login. The third principle assumes a breach has occurred; systems should be designed accordingly, using segmentation, encryption, endpoint detection, logging, and quick token revocation to limit the impact of a compromise.

NIST's logical model separates decision-making from enforcement, with the policy engine assessing risks, the administrator managing connections, and enforcement managing traffic [1]. This separation ensures policy consistency across users, devices, applications, and cloud workloads. Research shows that identity, endpoints, applications, data, networks, infrastructure, analytics, and automation should function as a unified system [3], [5]. However, literature reviews warn that much of the zero-trust evidence remains conceptual, with less focus on organisational and user impacts [4]. Therefore, ZTA isn't always better in every environment.

1.4 Technological Enablers

Identity governance is vital because zero trust decisions depend on identities. It includes verification, single sign-on, MFA, onboarding and offboarding procedures, regular reviews, duties segregation, and managing privileged access. NIST Guidelines suggest a lifecycle approach over passwords. Phishing-resistant methods like FIDO2

keys are safer than passwords as they link authentication to the legitimate service. Privileged access should be time-limited, approved, and logged. It also covers non-human identities such as service accounts and API keys, which may hold significant permissions.

ZTNA forces app level policies, unlike traditional VPNs that segment networks. It hides apps from unauthorised users and, after verification, limits connections, reducing discovery and lateral movement. However, it isn’t fully zero trust if reliance on outdated trust assumptions persists, especially if based on network location or broad access. Effective design includes per-app policies, device checks, session evaluation, and identity governance.

SASE combines networking and security through cloud points of presence, including SD-WAN, CASB, firewall-as-a-service, and ZTNA. It supports remote work by inspecting traffic near users, aiding in zero trust and connectivity modernisation. Its benefits are consistent policy and visibility, but challenges include provider dependency, data location, migration complexity, and vendor lock-in. SASE should enable, not define, architecture.

Micro-segmentation controls workload traffic with service mesh and mutual TLS. NIST applies zero trust to multi-cloud. Tools like CSPM, EDR, SIEM, and SOAR provide continuous monitoring, detecting unmanaged devices, risky logins, excessive permissions, or unusual data transfers, and alerting analysts or automating responses.

Figure 1. Context-aware Zero Trust decision cycle. [1], [5].




Table 1. Zero Trust technology enablers and limitations.

IGA, MFA and PAM	Assures identity, constrains privilege and supports rapid revocation.	Lifecycle errors, prompt fatigue and role explosion; require access reviews and break-glass tests.
ZTNA and SASE	Application-level access and consistent policy for hybrid users.	Provider concentration and opaque policy; require resilience, logging and exit provisions.
Micro-segmentation / service mesh	Contains lateral movement between workloads.	Discovery and rule complexity; baseline flows before staged enforcement.
CSPM, DLP and encryption	Reduces cloud misconfiguration and data exposure.	False positives and key dependency; assign ownership and recovery procedures.
SIEM/XDR and ATT&CK	Correlates risk signals and supports detection validation.	Telemetry cost and alert overload; tune to priority scenarios and response capacity.

1.5 Benefits, Challenges and Limitations

One advantage of ZTA is it reduces implied trust, preventing compromised users or workloads from accessing unrelated services. It enables remote work, multi-cloud use, third-party access, and improves audit evidence. Compliance controls are clearer since access, privileges, and actions are logged. Survivable ZTA research states that authenticating flows and dynamic policies increase resilience, but architecture must stay available during attacks [12]. Failure of identity providers or policy engines can halt business or cause unsafe access.

Implementation is challenging; organisations must understand users, devices, data flows, and dependencies. Older apps may lack current protocols for fine-grained authorisation. Poor documentation and mergers can cause duplicate directories and roles, risking disruption. The CISA maturity model views zero trust as progress in identity, devices, networks, apps, and data, supported by visibility, analytics, automation, and governance [7]. Cost is a hurdle, with licensing, integration, network redesign, skilled staff, and user support being costly. Research suggests benefits depend on organisation size, controls, and asset value [13]. Small startups may benefit more from identity consolidation and multi-factor authentication than from multiple platforms. Excessive verification can frustrate users, lead to workarounds, and cause false alarms. Privacy is crucial, as analytics and device monitoring gather detailed employee data; policies should be proportionate, transparent, and limit data retention.

Zero trust won't prevent all attacks. Authorised users can misuse data, applications have vulnerabilities, and sessions can be hijacked. Principles of ZTA depend on secure software, patches, incident response, and governance [14]. Gaps exist in interoperability, evaluation, scalability, and validation [15]. 'Never trust, always verify' means reducing inherited trust and

making risk decisions based on best evidence, not absolute certainty.

1.6 Real-world evidence and case studies

Google's BeyondCorp shows shifting access control from internal networks to user identity, device state, and application sensitivity, enabling VPN-free work [16]. It shows large-scale zero trust is feasible but requires years, accurate inventories, access proxies, device management, and adjustments for legacy workflows. It must be adapted to an organisation's resources and risk level before copying.

The 2024 Change Healthcare ransomware incident really teaches us an important lesson: attackers took advantage of compromised credentials to access a remote service that lacked MFA, which led to disruptions in healthcare payments. While having strong MFA is helpful, it's not entirely foolproof on its own. Combining extra measures like identity verification, limiting application access, segmentation, and quick revocation could have made it much harder for attackers to gain initial access or move laterally within the system. Threat data underscores the importance of layered security controls. For example, Verizon's 2026 DBIR reports that 31% of breaches started with exploited vulnerabilities, 48% involved ransomware, and third-party involvement has increased. Additionally, IBM's 2026 X-Force highlights how vital it is to protect public-facing applications and maintain strong identity and vulnerability protections. These ongoing trends remind us that Zero Trust Architecture (ZTA) needs to address both stolen identities and vulnerabilities in unpatched services to stay effective.

1.7 Recommended Implementation Approach

The implementation should proceed in phases based on risk assessment. First, an asset and data inventory must be carried out, key business services identified, and access routes mapped. Next, identity should be consolidated, inactive accounts deleted, phishing-resistant MFA introduced, and privileged access controlled via just-in-time methods. Then, ZTNA should be guided with a small user group and a limited set of high-value applications, gathering feedback and monitoring. Finally, segmentation, cloud posture management, and central telemetry should be added. Dynamic policies and automation should only be extended after establishing trust in data quality and operational processes.

Success should be measured by outcomes, not products deployed. Key metrics include protected critical applications, MFA coverage, removed dormant accounts, privileged sessions with time limits, closed high risk cloud findings, and recovery tests. Achieving zero trust maturity is gradual, and some controls may need recalling until systems improve.

1.8 Conclusion

Zero Trust Architecture is a reasonable answer to the challenges posed by cloud environments, remote work, and identity-based threats, since it eliminates trust based on location and ensures that each request is restricted to a specific resource openly authorised. Although identity governance, ZTNA, SASE, segmentation and continuous monitoring are all important enablers, none of them is adequate by itself. The main risks when implementing it are a lack of knowledge about assets, the presence of legacy systems, complex policies, cost, user inconvenience, privacy concerns and overreliance on vendors. There is evidence supporting a step-by-step program that starts with identity, critical assets, and visibility. Zero Trust should therefore be seen as a risk-management architecture that improves resilience, not as a solution that eliminates the need for secure development, patching, backups, incident response and governance.

TASK 2 - CISO PROPOSAL AND IMPLEMENTATION PLAN

The establishment of the Information Security Unit at Aegis FinTech Innovations Ltd.

2.1 Executive Summary

The 12-month proposal aims to establish an Information Security Unit at Aegis FinTech Innovations Ltd., a London startup developing an AI-driven financial platform and planning EU expansion. Based on Task 1, it highlights that identity compromise, accessible services, excessive privileges, and third-party access are key entry points for modern organisations. Zero trust requires identity, devices, apps, data, monitoring, and governance to develop concurrently. The strategy uses the NIST CSF 2.0 for governance and a phased Zero Trust architecture. Immediate goals include asset discovery, multi-factor authentication, access control, cloud management, third-party assurance, and ransomware recovery. A small team will handle risk and architecture, with a MSSP providing 24/7 monitoring until a larger security centre is feasible.

2.2 Work Stream 1: Organisational Design and Resourcing

The unit in question is headed by a CISO, who reports directly to the CEO and has access to the Board's Audit and Risk Committee. This ensures that security does not remain nothing more than an IT support role. During the first year, the GRC and Privacy function should have a Data Protection Officer (DPO) or arrange for an external DPO service. The DPO must be able to provide independent reporting on privacy risks. The Cloud Security Architect/Zero Trust Engineer will be responsible for designing the target architecture and for managing IAM, ZTNA and segmentation. The DevSecOps Engineer will incorporate security testing, secrets management and infrastructure-as-code checks into the product lifecycle. A Security Operations team, with support from an MSSP Response Analyst will handle alerts, investigations, and exercises. Initially, vendor risk should be managed within GRC rather than establishing a separate position.

The structure follows the eight zero trust dimensions from Task 1 [5], balancing independence and cost by handling governance and architecture internally while outsourcing high-volume monitoring with defined service levels. After 12 months, the Board should review incident volume and regulatory scope to decide if additional SOC analyst and IAM engineer are needed.




Figure 2. Proposed Information Security Unit. *The DPO retains independent access to the Board for privacy matters.

Role / initial FTE	Core accountability	Task 1 rationale
CISO — 1.0	Strategy, risk acceptance, Board reporting, regulatory and investor assurance.	F1: accountable governance.
GRC Lead — 1.0	Risk, policy, ISO/CSF evidence, vendor assurance and compliance coordination.	F1/F4: controls beyond access.

Cloud Security Architect — 1.0	Zero Trust patterns, IAM/PAM, CSPM, segmentation and resilience.	F2/F5: safe architecture and availability.
Security Operations Lead — 1.0	Detection engineering, incident command, threat intelligence and exercises.	F4: contain and respond.
DevSecOps Engineer — 1.0	Secure SDLC, dependency scanning, secrets and platform guardrails.	F4: vulnerable software remains a risk.
Security Analyst / IR — 1.0	Triage, investigations, tuning, vulnerability follow-up and metrics.	F3: measurable operations.
DPO / privacy counsel — 0.4 retained	Independent privacy advice, DPIAs, data-subject and regulator coordination.	F1: independent compliance challenge.

Table 3. Security unit roles and initial resourcing. Hiring can be phased, but accountability starts from month 1

2.3 Work Stream 2: the development of a policy framework

The policy states only approved AI and cloud services can handle company data, such as customer info, source code, credentials, and unreleased financial models, which must not be processed by public AI tools. Data should be classified before upload, and high-risk AI applications require privacy, security, and legal reviews. Employees must verify model outputs and not rely on generated content as financial advice. Approved services must use enterprise accounts, encryption, logging, least privilege, and clear data retention policies. The policy also requires reporting shadow AI, prompt injection, disclosures, and unauthorised integrations.

Policy requires workforce access through the company's identity provider, using phishing-resistant multi-factor authentication and zero-trust network access, preventing direct exposure to management interfaces. Access is limited to enrolled devices that meet standards for patches, encryption, screen locking, and endpoint detection. Administrators must use separate privileged accounts and obtain just-in-time privileges via PAM. Access is specific to necessary applications, with high-risk sessions re-authenticated or blocked. Local admin rights, shared accounts, and unmanaged storage are forbidden unless authorised. This policy follows Task 1's conclusion that ZTNA must incorporate both identity and device context, not just mimic a VPN.

Policy on managing third-party and cloud vendor risks. Before contracts, all suppliers must be assigned a risk level. High-risk vendors must provide evidence like ISO 27001, SOC 2 reports, penetration test summaries, data flow info, and sub-processor lists. Contracts should include clauses on encryption, access control, breach notification, data location, secure deletion, audit rights, business continuity, exit support, and sub-contracting restrictions. For DORA-related services, a register of ICT providers and contracts is required. Concentration risk must be managed to avoid single points of failure. Continuous monitoring and annual reassessment are essential, not just initial questionnaires.



Policy	Executive summary and mandatory controls	Owner / evidence
Acceptable Use of AI and Cloud	Only approved services and classified data may be used. Personal, customer, or source data must not enter public generative-AI tools. Require human validation of outputs, model/data provenance, secure prompts/secrets, logging, retention limits, and copyright/privacy review.	CTO; service register, DLP events, model cards, training and exception log.
Hybrid Work and Remote Access	Managed and encrypted devices, phishing-resistant MFA, ZTNA per application, device compliance, least privilege, restricted local admin, and secure collaboration are mandatory. Privileged work needs hardened workstations and recorded PAM sessions.	COO/CISO; device compliance, ZTNA coverage, access reviews and remote-access alerts.

Third-Party and Cloud Vendor Risk	Tier vendors by data, privilege, criticality, and substitutability. Due diligence includes security, privacy, resilience, subcontractors, and location. Contracts specify incident notice, audit evidence, deletion/return, recovery, change control, and exit support.	Procurement/GRC; inventory, risk decisions, contracts, attestations and exit tests.
--	---	---

Table 4. Executive summaries of the three mandatory policies.

2.4 Work Stream 3: a regulatory compliance strategy

Aegis must embed compliance into product and operational design. Since 17 January 2025, the EU's GDPR has applied to personal data processing, requiring lawful, fair, and transparent handling, data minimisation, privacy by design, security measures, record-keeping, breach procedures, and Data Protection Impact Assessments for high-risk processing [21]. The company should determine whether it needs an EU representative and identify the lead supervisory authority. DORA has been in force since 17 January 2025, covering ICT risk management, incident handling, resilience testing, and third-party risks for regulated financial entities [22], [23]. A legal scoping exercise is needed, as not all tech startups qualify as financial entities, but Aegis may fall within DORA's scope as an ICT supplier to regulated clients.

Phase	GDPR deliverables	DORA / resilience deliverables	Assurance gate
0–90 days	Data map, processing records, lawful-basis review, retention baseline, DPO assessment.	Scope/legal applicability, critical service map, ICT asset/vendor registers and risk tolerance.	Board approves scope, owners and priority gaps.
Months 4–6	DPIA method, rights/breach procedures, processor clauses, privacy-by-design checkpoints.	Incident taxonomy, response roles, resilience test plan and vendor contract schedule.	Evidence sample and first tabletop.
Months 7–9	Remediate high-risk processing; test deletion, access and notification workflows.	Test critical scenarios, recovery and key vendors; record and own findings.	Steering Committee accepts residual risk.
Months 10–12	Internal audit, metrics, training refresh and management review.	DORA-aligned evidence pack, lessons learned and next-year testing plan.	Board assurance statement and investor pack.

Table 5. GDPR and DORA compliance roadmap.

2.5 Work Stream 4: technical standards and architecture

NIST CSF 2.0 is recommended as the main security framework because it is technology-neutral, suitable for all sizes, and based on Govern, Identify, Protect, Detect, Respond, and Recover steps [20]. The updated Govern function helps a Series B company connect cyber risk with Board oversight, supplier management, legal obligations, and business priorities. While ISO 27001:2022 can be considered later for certification, starting with NIST CSF allows Aegis to develop current and target profiles and an improvement plan without focusing on certification in the first year.

The architecture places the identity provider, IGA, phishing-resistant MFA, PAM, and device positioning before a zero-trust network access (ZTNA) or Secure Access Service Edge (SASE) layer. The policy engine assesses user or workload access based on their identity, device, risk level, and data sensitivity. Public APIs are protected by an API gateway, web application firewall, service identity, and rate limiting. Cloud workloads are separated with encrypted traffic between services. Cloud Security Posture Management (CSPM) and Cloud Native Application Protection Platform (CNAPP) detect misconfigurations, while Endpoint Detection and Response (EDR) and Security Information and Event Management (SIEM) with Security Orchestration, Automation, and Response (SOAR) provide continuous monitoring. Data protection includes loss prevention, key management, and encryption, supported by immutable backups for recovery. This design demonstrates that zero trust requires integrating identity, applications, workloads, data, and telemetry, instead of relying solely on a single gateway.

Figure 3. Zero Trust architecture for Aegis FinTech from Task 1 and NIST SP 800-207/207A [1], [11].

Users & service identities →	Identity and device trust →	Policy decision →	Enforcement points →	Protected resources
Employees Third parties Workloads	IdP / IGA / PAM EDR / MDM Risk signals	Policy engine Least privilege Step-up MFA	ZTNA / API gateway micro-segmentation Service mesh	AI analytics platform Data stores SaaS and cloud
Continuous telemetry and assurance: SIEM/XDR • CSPM • vulnerability management • DLP • threat intelligence • immutable audit and recovery evidence				

2.6 Work Stream 5: strategic risk assessment

The risk assessment uses a 1-5 scale for likelihood and impact. Scores 15-25 are high, 8-14 medium, and 1-7 low. Main risks come from Task 1 and the 2026 DBIR, especially regarding vulnerability exploitation, ransomware, and third-party involvement [18]. The security team reviews the risk register monthly and the Steering Committee quarterly, with owners and target dates, treating it as a dynamic document.

Risk / inherent score	Key treatment	Owner / target residual / indicator
Credential or cloud-control-plane compromise — 20 (5×4)	Phishing-resistant MFA, PAM/JIT, device trust, CSPM, separation of duties, anomaly detection and emergency access tests.	Cloud Security Architect; 8. KRIs: standing admins, unmanaged sessions, critical misconfigurations.
Third-party or software supply-chain compromise — 20 (4×5)	Tiered due diligence, SBOM/dependency scanning, signed releases, restricted vendor access, contract notification, segmentation and exit testing.	GRC Lead/CTO; 10. KRIs: unassessed critical vendors, critical dependency age, overdue clauses.
Ransomware causing data/analytics outage — 20 (4×5)	EDR/XDR, segmentation, patching, tested immutable backups, response retainer, crisis communications and recovery exercises.	Security Operations Lead; 8. KRIs: backup test failures, unpatched critical exposure, recovery variance.

Table 6. High-level strategic risk register (reviewed regularly).

2.7 Work Stream 6: incident response and business continuity

The ransomware response procedure has six stages. In preparation, maintain an incident contact list, clean system images, keep offline emergency accounts, have endpoint tools ready, prepare network maps, and create ransomware-resistant backups. For detection and triage, open an incident record, identify affected people, devices, and services, preserve logs, and determine if data was stolen. During containment, revoke tokens, disable compromised accounts, isolate devices and cloud workloads, block indicators, and use ZTNA policies and segmentation to prevent lateral movement. In eradication, remove persistence, patch vulnerabilities, rotate passwords, keys, and secrets, and rebuild systems from trusted images. During recovery, restore services based on priorities using immutable backups, verify data integrity, increase monitoring before reconnecting. Finally, conduct a lessons learned review, update controls, and report metrics to the Board

Prepare →	Detect & declare →	Contain & preserve →	Eradicate & restore →	Learn & assure
---------------------	----------------------------------	------------------------------------	-------------------------------------	---------------------------

Figure 4. Ransomware response and recovery lifecycle. Legal, privacy, communications, and evidence preservation operate throughout [25], [26].

Communications must proceed in parallel. The CISO takes incident command, with the Legal department and DPO handling regulatory notification, the CEO approving external statements, and the COO managing customer service. In case of a reportable personal data breach, notification to the regulator must be prompt, ideally within 72 hours [27]. Incident reporting under DORA requires following specific templates and timeframes. The company should also contact its insurer, NCSC, law enforcement, and key suppliers. Paying ransom is generally not advised, as it doesn't guarantee data recovery or deletion and requires legal, sanctions,

ethical, and Board review. NCSC advises resilient, unalterable backups [25], and CISA recommends a tested recovery process and documented response plan [26].

2.8 Work Stream 7: governance and the 12-month strategic roadmap

The Cyber Security Steering Committee, chaired by the CEO, includes the CISO, CTO, CFO, COO, Head of Legal/{DPO}, Head of Product/AI, HR lead, and an independent adviser or internal audit representative. It meets monthly for six months and quarterly thereafter. Its roles are to approve risk appetite, review the risk register, resolve ownership conflicts, monitor regulatory readiness, oversee key suppliers, and accept high risks. The CISO provides a one-page dashboard on control coverage, incidents, overdue actions, supplier risk, and recovery readiness.

Member	Decision contribution	Cadence / escalation
CEO (Chair), CISO, CTO, COO, CFO	Risk appetite, funding, architecture and operational priorities.	Monthly; severe incidents or tolerance breach within 24 hours.
Legal/DPO, Product, Data, HR, Procurement	Regulation, customer impact, people, data and third-party ownership.	Monthly; join incidents and material change reviews.
Board Risk Committee	Challenges management, accepts material residual risk and oversees assurance.	Quarterly and on critical escalation.

Table 7. Cyber Security Steering Committee.

In the first quarter, the roadmap emphasises identity and public exposure, as Task 1 and recent threat data show that valid credentials and vulnerable services can cause serious incidents. Later quarters focus on scaling segmentation, monitoring, and assurance. Key performance indicators include 100% MFA use for workforce and admins, 100% privileged account coverage by PAM, at least 95% device compliance, monitoring of all critical cloud accounts by CSPM, resolving critical vulnerabilities within SLA, assessing high-risk vendors before access, quarterly access reviews, backup and restore success above 95%, and completing incident exercises with all actions closed. Metrics should guide decisions, not suggest security is complete.

Period	Priority deliverables	Board-level KPIs / target
Q1: Govern and discover	Hire leads; establish committee; inventories; data map; critical services; baseline policies and incident contacts.	≥95% critical assets with owner; 100% privileged accounts inventoried; critical vendors tiered.
Q2: Control priority paths	IGA lifecycle, phishing-resistant MFA/PAM pilot, ZTNA for admins/vendors, CSPM, policy approval, first tabletop.	100% privileged users phishing-resistant MFA; median leaver revocation <4 hours; ≥90% critical cloud controls compliant.
Q3: Contain and prove	Workload identity, segmentation, SIEM/XDR use cases, vendor remediation, DPIA workflow, recovery exercise.	≥90% critical applications behind ZTNA; ≥85% priority ATT&CK detection coverage tested; recovery within approved RTO.
Q4: Assure and scale	Independent review, immutable restore test, DORA/GDPR evidence pack, roadmap refresh and Series B assurance.	100% critical restores tested; high-risk audit findings owned within 10 days; exceptions past expiry = 0.

Table 8. Twelve-month implementation roadmap and KPIs.

2.9 Conclusion

Aegis can build a trustworthy security setup within 12 months by combining clear ownership, specialised internal knowledge, and managed monitoring. The NIST CSF 2.0 provides a helpful governance framework, and the Zero Trust Architecture turns the insights from Task 1 into practical controls for identity, devices, applications, cloud workloads, and data. The plan thoughtfully starts with the basics and high-risk services, avoiding an expensive, full-scale overhaul all at once. Success will rely on the Board's involvement, having reliable inventories, practicing disciplined identity governance, managing suppliers carefully, testing recovery procedures, and continuously improving as the company grows into the EU.

REFERENCES

- [1] S. Rose, O. Borchert, S. Mitchell and S. Connelly, “Zero Trust Architecture,” NIST Special Publication 800-207, Aug. 2020. doi: 10.6028/NIST.SP.800-207. [Available online](#) (accessed Jul. 7, 2026).
- [2] UK National Cyber Security Centre, “Introduction to Zero Trust,” Jul. 2021. [Available online](#) (accessed Jul. 7, 2026).
- [3] N. F. Syed, S. W. Shah, A. Shaghaghi, A. Anwar, Z. Baig and R. Doss, “Zero Trust Architecture (ZTA): A Comprehensive Survey,” IEEE Access, vol. 10, pp. 57143-57179, 2022. doi: 10.1109/ACCESS.2022.3174679. [Available online](#) (accessed Jul. 7, 2026).
- [4] 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. 102436, 2021. doi: 10.1016/j.cose.2021.102436. [Available online](#) (accessed Jul. 7, 2026).
- [5] 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. 103412, 2023. doi: 10.1016/j.cose.2023.103412. [Available online](#) (accessed Jul. 7, 2026).
- [6] D. Temoshok et al., “Digital Identity Guidelines,” NIST Special Publication 800-63-4, Jul. 2025. doi: 10.6028/NIST.SP.800-63-4. [Available online](#) (accessed Jul. 7, 2026).
- [7] Cybersecurity and Infrastructure Security Agency, “Zero Trust Maturity Model, Version 2.0,” Apr. 2023. [Available online](#) (accessed Jul. 7, 2026).
- [8] UK National Cyber Security Centre, “Authenticate and authorise everywhere,” Jul. 2021. [Available online](#) (accessed Jul. 7, 2026).
- [9] Cybersecurity and Infrastructure Security Agency, “The Journey to Zero Trust: Using SASE in a Modern TIC 3.0 Solution,” Jul. 2026. [Available online](#) (accessed Jul. 7, 2026).
- [10] UK National Cyber Security Centre, “Designing secure access with ZTNA,” May 2026. [Available online](#) (accessed Jul. 7, 2026).
- [11] 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, Sep. 2023. doi: 10.6028/NIST.SP.800-207A. [Available online](#) (accessed Jul. 7, 2026).
- [12] L. Ferretti, F. Magnanini, M. Andreolini and M. Colajanni, “Survivable zero trust for cloud computing environments,” Computers & Security, vol. 110, art. 102419, 2021. doi: 10.1016/j.cose.2021.102419. [Available online](#) (accessed Jul. 7, 2026).
- [13] Z. Adahman, A. W. Malik and Z. Anwar, “An analysis of zero-trust architecture and its cost-effectiveness for organisational security,” Computers & Security, vol. 122, art. 102911, 2022. doi: 10.1016/j.cose.2022.102911. [Available online](#) (accessed Jul. 7, 2026).
- [14] E. B. Fernandez and A. Brazhuk, “A critical analysis of Zero Trust Architecture (ZTA),” Computer Standards & Interfaces, vol. 89, art. 103832, 2024. doi: 10.1016/j.csi.2024.103832. [Available online](#) (accessed Jul. 7, 2026).
- [15] S. Mushtaq, M. Mohsin and M. M. Mushtaq, “A Systematic Literature Review on the Implementation and Challenges of Zero Trust Architecture Across Domains,” Sensors, vol. 25, no. 19, art. 6118, 2025. doi: 10.3390/s25196118. [Available online](#) (accessed Jul. 7, 2026).
- [16] Google Cloud, “BeyondCorp Zero Trust Enterprise Security.” [Available online](#) (accessed Jul. 7, 2026).
- [17] U.S. House Committee on Energy and Commerce, “What We Learned: Change Healthcare Cyber Attack,” May 2024. [Available online](#) (accessed Jul. 7, 2026).
- [18] Verizon, “2026 Data Breach Investigations Report,” 2026. [Available online](#) (accessed Jul. 7, 2026).
- [19] IBM Security, “X-Force Threat Intelligence Index 2026,” 2026. [Available online](#) (accessed Jul. 7, 2026).
- [20] National Institute of Standards and Technology, “The NIST Cybersecurity Framework (CSF) 2.0,” Feb. 2024. doi: 10.6028/NIST.CSWP.29. [Available online](#) (accessed Jul. 7, 2026).
- [21] European Parliament and Council, “Regulation (EU) 2016/679 (General Data Protection Regulation),” Apr. 2016. [Available online](#) (accessed Jul. 7, 2026).

- [22] European Parliament and Council, “Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA),” Dec. 2022. [Available online](#) (accessed Jul. 7, 2026).
- [23] European Insurance and Occupational Pensions Authority, “Digital Operational Resilience Act (DORA),” 2025. [Available online](#) (accessed Jul. 7, 2026).
- [24] Cybersecurity and Infrastructure Security Agency, “Implementing Phishing-Resistant MFA.” [Available online](#) (accessed Jul. 7, 2026).
- [25] UK National Cyber Security Centre, “Principles for ransomware-resistant cloud backups,” Oct. 2023. [Available online](#) (accessed Jul. 7, 2026).
- [26] Cybersecurity and Infrastructure Security Agency, “#StopRansomware Guide,” 2025. [Available online](#) (accessed Jul. 7, 2026).
- [27] Information Commissioner’s Office, “72 hours - how to respond to a personal data breach.” [Available online](#) (accessed Jul. 7, 2026).