

RESEARCH AND CASE SCENARIO IN DATA SECURITY AND COMPLIANCE

ACCA5034



AUGUST 6, 2026
CRISTINEL-COSMIN VANUSCA
UWTSD LONDON

Content

TASK 1.....	2
Introduction.....	3
1. Traditional Perimeter-Based Security.....	3
2. Zero Trust Architecture: Definition and Core Principles.....	4
3. How Zero Trust Helps Organisations Anticipate Threats.....	4
4. Protecting Information Assets.....	5
5. Detecting Threats Through Continuous Monitoring.....	6
6. Automated Response and Threat Containment.....	7
7. Zero Trust Compared with Traditional Security.....	7
8. Implementation Challenges.....	8
9. Practical Implementation Approach.....	9
TASK 2.....	11
Introduction.....	11
1. Organisational Structure and Accountability.....	12
2. Policies and Procedures.....	14
3. Governance and Strategic Alignment.....	15
4. Risk Management and Assessment.....	16
5. Compliance, Incident Response and Protection Controls.....	19
6. Recommended Cybersecurity Programme Diagram.....	20
7. Performance Measurement.....	21
Overall Conclusion.....	22
References.....	23

Assignment Title:	Research and case scenario in Data Security and Compliance	Word count (or equivalent):	5000 words or equivalent
-------------------	--	-----------------------------	--------------------------

TASK 1 – RESEARCH IN DATA SECURITY AND COMPLIANCE

The Zero Trust Architecture

A Critical Analysis of Its Principles, Implementation Challenges and Efficacy Compared with Traditional Models

From lecturer:


Task 1 – Research in Data Security and Compliance

Conduct research on one of the following contemporary challenge areas below. Your paper must analyse and critically evaluate how modern organizations can create an environment that anticipates, protects, detects, and automatically responds to threats.

Your research paper must not be more or less than 10% of 2500 words, choose one or more of the following:

- Securing the Hybrid Workforce: Protecting data beyond the corporate perimeter (ZTNA, SASE, DLP).
- Governance of AI and Machine Learning Systems: Managing risk, ethics, and compliance for AI data and models.
- The Zero Trust Architecture: A critical analysis of its principles, implementation challenges, and efficacy compared to traditional models.
- Compliance in a Fragmented World: Navigating GDPR, alongside new regulations like DORA (Digital Operational Resilience Act) and NIS2 for critical infrastructure.

Should you want to choose a different title, you need to obtain the consent from the lecturers; otherwise, it may lead to the failure of the assignment.



Research topic: The Zero Trust Architecture: A Critical Analysis of Its Principles, Implementation Challenges, and Efficacy Compared with Traditional Security Models.

Introduction

Modern organisations rely on cloud services, hybrid work, mobile devices and third-party services. Users access distributed data and applications from changing locations, networks and devices, weakening reliance on a fixed organisational boundary [1]. Traditional perimeter security assumes that internal users and devices are trusted while external entities are threats. If an attacker bypasses the boundary or obtains valid credentials, broad internal access may permit lateral movement with limited resistance [1]. Zero Trust Architecture removes implicit trust based on network location. Access is evaluated using identity, device condition, resource sensitivity and contextual risk [4]. This research evaluates whether Zero Trust helps organisations anticipate, protect against, detect and automatically respond to threats more effectively than perimeter-based security. It argues that continuous verification, least privilege and segmentation improve protection, but success depends on asset visibility, identity governance, integration, organisational maturity and sustained investment [3].

1. Traditional Perimeter-Based Security


Traditional perimeter security follows a ‘castle-and-moat’ model: firewalls protect the network boundary and VPNs provide remote connectivity. Its relative simplicity, familiar administration and lower initial cost explain its continued use.

Its main weakness is implicit internal trust. After compromising credentials or a device, an attacker may enter through legitimate channels, escalate privileges and move laterally, insiders may also possess excessive access [6].

Cloud platforms, mobile devices, remote work and third-party integrations have weakened the concept of a fixed perimeter because data and applications increasingly operate outside the organisation’s direct network control [3].

Perimeter controls remain valuable layers, but they cannot independently secure a distributed organisation. They should be combined with identity-based access, least privilege, monitoring and granular segmentation [1].

Figure 1. Traditional perimeter-based “castle-and-moat” model.



2. Zero Trust Architecture: Definition and Core Principles


Zero Trust Architecture removes implicit trust from access decisions. ‘Never trust, always verify’ means that access is based on current evidence, including identity, device posture, resource sensitivity and contextual risk, rather than network location , [4]. The ‘assume breach’ principle treats accounts, devices or network components as potentially compromised. Least-privilege controls restrict lateral movement by granting only the permissions required for a specific task . [2].

Access is authorised per session and resource. Decisions can consider role, authentication strength, location, behaviour and device health, while micro_segmentation and continuous monitoring allow access to be reassessed as conditions change [1].

NIST identifies three logical components: the Policy Engine decides whether access is allowed, the Policy Administrator creates or terminates the communication path, and the Policy Enforcement Point applies that decision [1].

Zero Trust is not a single product or the elimination of all trust. It is dynamic trust management in which limited access is established from evidence, reassessed continuously and withdrawn when risk changes [3].

Figure 2. Core logical components of a Zero Trust Architecture.




3. How Zero Trust Helps Organisations Anticipate Threats

Zero Trust improves threat anticipation through visibility across users, devices, applications, workloads and data. Accurate inventories reveal authorised assets, unmanaged systems, inactive identities and excessive permissions before exploitation [2]. Vulnerability scanning identifies known weaknesses and insecure configurations. Combined with threat intelligence, it supports risk scoring based on exploit activity, exposure, identity, device health and resource sensitivity [9]. User and Entity Behaviour

Analytics can detect abnormal logins, data access and privilege changes. Device-posture checks assess management, patching and encryption, while data classification applies stronger controls to sensitive information, [6]. Attempt to access sensitive data from an unmanaged or unpatched device may therefore trigger additional authentication, restricted access or blocking before the activity becomes a serious incident. This capability depends on data quality. Incomplete inventories, inaccurate identities or missing telemetry can produce poor decisions, so continuous discovery and validation are essential [2].

Figure 3. Contextual Zero Trust access decision.




4. Protecting Information Assets

Zero Trust protects assets through identity controls, device security, data protection and segmentation. Identity and Access Management defines who may access each resource, while Multi-Factor Authentication reduces reliance on passwords [7]. Privileged Access Management limits and monitors administrator and service accounts, supporting least privilege by granting elevated access only when required [6]. Role-Based Access Control links permissions to job responsibilities, whereas Attribute-Based Access Control also considers department, location, device posture and data sensitivity. Micro_segmentation restricts movement between systems [4]. Encryption protects data at rest and in transit, endpoint tools detect malware and insecure configurations. Data Loss Prevention restricts unauthorised transfers, and workload controls validate service identities across cloud and on-premises systems [6].

For remote users, Zero Trust Network Access exposes specific applications instead of an entire internal network and can reassess access using identity, device posture and risk, [3].

These controls reduce lateral movement and the impact of compromised accounts, but they also increase complexity. Effective deployment requires usable policies, reliable integration and a balance between protection and productivity.

Figure 4. Layered protection of information assets in a Zero Trust environment.




5. Detecting Threats Through Continuous Monitoring

Zero Trust supports detection through continuous monitoring of identities, endpoints, applications, networks and cloud services. SIEM platforms correlate these records to identify patterns that isolated tools may miss [5]. Endpoint Detection and Response provides detailed visibility into processes, malware behaviour and unauthorised changes. Identity logs can reveal unusual locations, repeated failures or abnormal login times. User and Entity Behaviour Analytics can establish normal patterns and identify deviations associated with compromised accounts or insider activity. Context such as device health, current privilege and resource sensitivity helps prioritise suspicious downloads, unrelated access and communication between segmented environments [6].

More data does not automatically create better detection. Poor configuration can generate false positives, alert fatigue and missed incidents. Detection can also be weakened by incomplete logging, limited cloud visibility, skills shortages and integration problems. Effective monitoring therefore depends on reliable data sources, tuned analytics, tested integrations and sufficient operational capacity [5].

Figure 5. Continuous monitoring and detection pipeline.



6. Automated Response and Threat Containment

Zero Trust can improve incident response by combining continuous monitoring with automated policy enforcement. Security Orchestration, Automation and Response platforms integrate alerts from SIEM, EDR, identity systems and cloud services, then apply predefined playbooks to coordinate containment [8]. Automated actions may include revoking access tokens, forcing reauthentication, suspending accounts, isolating endpoints, blocking network connections, reducing privileges, preventing sensitive downloads and creating incident tickets. Automation reduces manual delay and enables containment to begin when suspicious activity is detected. Response actions should be divided into categories. Fully automatic actions are suitable for high-confidence, low-disruption situations. Automatic but reversible actions can temporarily isolate a device or restrict access while allowing rapid restoration. More disruptive actions, such as disabling critical accounts or shutting down production services, should normally require analyst approval [8]. Automation can reduce attacker dwell time and organisational response time. However, incorrect rules, incomplete context and false positives may disrupt legitimate operations. Automated response should therefore rely on tested playbooks, clear escalation procedures, audit trails and proportionate human oversight [8].




Figure 6. Levels of automated response and required oversight.

7. Zero Trust Compared with Traditional Security

Traditional perimeter-based security and Zero Trust Architecture differ primarily in the basis of trust, access granularity, authentication, segmentation, monitoring and response. Traditional security relies heavily on network location, while Zero Trust evaluates identity, device condition and context for each protected resource [2]. Zero Trust is more suitable for cloud services, mobile devices and hybrid working environments. Traditional controls remain useful as part of a multilayered strategy, but they should be combined with identity-based access, segmentation, continuous monitoring and automated response [3].

Tab.1

Security criterion	Traditional security	Zero Trust Architecture
Basis of trust	Network location	Identity, device posture and context
Access	Usually broad after network entry	Granular and resource-specific
Authentication	Predominantly at connection	Continuous and contextual
Segmentation	Large or limited network zones	Micro_segmentation
Remote access	VPN-based network access	ZTNA application access
Monitoring	Often fragmented	Continuous and integrated
Response	Frequently manual	Can be automated
Threat model	Attacker assumed mainly outside	Compromise is assumed

8. Implementation Challenges

Although Zero Trust offers significant security benefits, implementation can be technically, financially and organisationally difficult. Legacy systems may not support modern identity controls, continuous authentication or granular policies. Incomplete asset inventories and shadow IT create gaps outside formal monitoring and enforcement [3]. Service accounts may operate automatically, possess excessive privileges and lack support for interactive MFA. Integration complexity can create inconsistent policy decisions across identity, endpoint, network and cloud systems. Organisations may also lack personnel with the skills required to design policies, integrate tools and analyse security data.

Tab.2


Challenge	Potential impact
Legacy systems	Limited compatibility with modern controls
Incomplete inventories	Unprotected or unmanaged assets
Integration complexity	Inconsistent policies and security gaps
Policy misconfiguration	Legitimate access blocked or excessive access granted
Identity-system failure	Widespread loss of access to resources
Excessive authentication	Reduced productivity and user frustration
Vendor lock-in	Reduced flexibility and higher migration cost

Challenge	Potential impact
Continuous monitoring	Privacy concerns and increased governance requirements

9. Practical Implementation Approach

Zero Trust should be implemented gradually rather than introduced as a single organisation-wide transformation. A phased approach reduces disruption and allows security teams to test policies, resolve integration problems and demonstrate measurable improvements [3]. The first stage is to identify critical data, applications and systems. The organisation should then create an accurate inventory of users, devices, workloads, applications and service accounts. Identity management should be strengthened by removing inactive accounts, defining ownership and introducing MFA. Excessive permissions should be removed, and critical systems should be segmented to restrict unnecessary communication. Endpoint, network, identity and cloud telemetry should then be integrated into central monitoring. Automated response should be introduced gradually, beginning with low-risk and reversible actions. More disruptive actions should initially require analyst approval. Policies must be reviewed continuously as systems, threats and responsibilities change [5].

Figure 7. Gradual Zero Trust implementation process.



Tab.3

Performance indicator	Purpose
Number of privileged accounts	Measures privilege reduction
MFA coverage	Shows identity-control adoption

Performance indicator	Purpose
Percentage of managed devices	Measures endpoint visibility
Mean time to detect	Evaluates detection speed
Mean time to respond	Evaluates containment speed
Number of policy exceptions	Identifies control weaknesses
Lateral movement incidents	Measures segmentation effectiveness
False-positive rate	Evaluates monitoring accuracy

TASK 2 – CASE SCENARIO

The Equifax Data Breach


A Structured Cybersecurity Programme Using NIST CSF 2.0 and Zero Trust

From lecturer :

Task 2 – Case Scenario

The 2017 Equifax data breach provides a critical case study for designing a structured, forward looking cybersecurity programme. The incident exposed the personal data of approximately 147 million individuals following the exploitation of a known Apache Struts vulnerability that had not been patched. Attackers maintained access for over two months, extracting highly sensitive information including Social Security numbers, dates of birth, and addresses. The breach reflects not only technical failure but also weaknesses in governance, accountability, risk oversight, and compliance management.

Using this case as the foundation, students are required to apply the knowledge and understanding developed in **Task 1** and lecture content to design a structured cybersecurity programme. Responses must demonstrate analytical reasoning, application of recognised frameworks, and justified decision making rather than simple description of the incident.



Introduction


The 2017 Equifax breach shows how a technical weakness can become a major organisational and regulatory failure. Attackers exploited an unpatched Apache Struts vulnerability in the online dispute portal, remained active for about 76 days and accessed sensitive information relating to approximately 147 million people [14]. The incident was not caused only by one missing patch. Investigations identified weaknesses in asset visibility, vulnerability ownership, segmentation, monitoring, governance and escalation. The FTC also alleged that remediation was not verified and that an expired certificate prevented inspection of encrypted traffic [17].

This task proposes a forward-looking programme structured around the NIST CSF 2.0 functions of Govern, Identify, Protect, Detect, Respond and Recover [5]. It also applies

Zero Trust principles so that access is continually evaluated rather than trusted because of network location [1].

Effective prevention requires executive oversight, measurable policies, risk-based controls, clear ownership, verified remediation and continuous monitoring; technology cannot compensate for weak accountability.

Figure 8. Key timeline of the 2017 Equifax breach.



1. Organisational Structure and Accountability

1.1 Board-level accountability

Cybersecurity should be governed as an enterprise risk. The Board, supported by a competent Risk or Cybersecurity Committee, should approve risk appetite, oversee investment and receive reports on critical vulnerabilities, overdue remediation, incidents, control failures and regulatory exposure [6].


Controls and funding should be proportionate to data sensitivity and business impact, consistent with the CSF 2.0 Govern function [5].

1.2 Executive security leadership

A Chief Information Security Officer should lead the programme, report to the Chief Executive Officer and have direct access to the Board Risk Committee. The CISO should control security policy, architecture, risk, vulnerability governance and incident response, while the CIO remains responsible for technology delivery.

This separation limits conflicts between security and operational pressures. Equifax demonstrated how fragmented ownership and weak verification can leave a known vulnerability exploitable [15].

Figure 9. Proposed cybersecurity governance structure.



1.3 Operational roles

Security Architecture should define design standards and approve high-risk changes. Vulnerability and Patch Management should identify, prioritise and verify remediation. The Security Operations Centre should monitor alerts, while Incident Response coordinates containment and recovery. DevSecOps should embed security testing and dependency management into development. Identity and Access Management should control authentication, privileged access and access reviews. The Data Protection Officer should oversee privacy and notification duties. Risk and Compliance should map controls to legal and contractual obligations. Business Continuity and Disaster Recovery should maintain resilience, while Internal Audit independently tests control effectiveness [13].

1.4 Accountability model

Tab.4

Activity	Accountable	Responsible	Consulted / Informed
Critical patch approval	CISO	Vulnerability Management Team	CIO; asset owner
Patch deployment	CIO	IT Operations	CISO; business owner
Secure development	Head of Engineering	DevSecOps Team	Security Architecture
Incident declaration	CISO	Incident Response Manager	CEO; Legal; DPO

Activity	Accountable	Responsible	Consulted / Informed
Regulatory notification	Data Protection Officer	Legal and Compliance	CISO; senior management
Board reporting	CISO	Cyber Risk Team	CEO; Internal Audit

Roles require formal authority, resources, measurable objectives and escalation thresholds. Decisions, exceptions and remediation deadlines must be documented and reviewed.

2. Policies and Procedures

Cybersecurity policies must be approved, enforceable, measurable and auditable. Each should define ownership, responsibilities, evidence, metrics, exceptions and consequences. Equifax shows that instructions are ineffective unless remediation is verified and unresolved risks are escalated [15].

2.1 Vulnerability management policy

Continuous scanning should cover endpoints, servers, cloud resources, applications and internet-facing systems. Findings should be prioritised by severity, exploitability, exposure, data sensitivity and business impact, assigned to a named owner and tracked centrally [12].


Critical vulnerabilities should normally be remediated within 48–72 hours and high-risk findings within 7–14 days. Exceptions require justification, compensating controls, an expiry date and CISO approval; overdue critical items should be escalated.

2.2 Patch management policy

Security teams should monitor vendor advisories and threat intelligence. Emergency procedures should support accelerated testing and deployment, with rollback arrangements for operational failure.

Patch records must prove successful installation through automated verification, rescanning and audit. This directly addresses Equifax's failure to confirm that the Apache Struts patch had been applied [17].

Figure 10. Vulnerability and patch-management lifecycle.



2.3 Secure Software Development Lifecycle policy

Security requirements should be built into design through threat modelling, secure coding, dependency scanning, SAST, DAST, peer review and penetration testing. Production release should require documented security approval [11].

2.4 Asset visibility policy

Accurate inventories should cover hardware, software, cloud services and applications, with an owner, purpose and classification for each asset. Automated discovery should identify unmanaged or internet-facing systems, while obsolete assets should be retired or isolated [5].

2.5 Monitoring and logging policy

Security logs should be centralised in SIEM, protected from alteration and governed by defined retention, alert, investigation and escalation requirements [6].

2.6 Data protection policy

Sensitive information should be classified, encrypted and restricted by least privilege. Data minimisation, retention schedules, secure deletion, DLP and periodic access reviews should reduce exposure [6].

Tab.5

Policy area	Target or evidence
Critical vulnerabilities	Remediation within 48–72 hours
High-risk vulnerabilities	Remediation within 7–14 days
Patch deployment	Installation evidence and successful rescan
Asset inventory	Named owner and classification for every critical asset
Monitoring	Centralised, complete and protected logs
Access control	Periodic access-review records
Policy exceptions	Approved owner, compensating control and expiry date

Implementation should be measured through remediation compliance, patch success, inventory coverage, exceptions and audit findings, with repeated non-compliance escalated for corrective action.

3. Governance and Strategic Alignment

The programme follows the six NIST CSF 2.0 functions—Govern, Identify, Protect, Detect, Respond and Recover—to organise and communicate cybersecurity outcomes [5].

3.1 Govern

Govern establishes direction and accountability. The Board should approve strategy and risk appetite, assign ownership, address legal and supply-chain obligations, and review investment, exceptions and major residual risks [6].

3.2 Identify

Identify requires current inventories of assets, identities and data, supported by business-impact analysis, classification and risk assessment. Systems holding identity and financial information should receive the highest priority [9].

3.3 Protect

Protect should combine least privilege, strong authentication, encryption, secure configuration, awareness and verified patching. Internet-facing applications and sensitive databases require stricter segmentation and change control [12].

3.4 Detect

Detect should correlate SIEM, EDR, vulnerability and anomaly data across identity, endpoint, network, application and cloud sources to identify unusual access, unpatched systems and abnormal extraction [6].


3.5 Respond

Respond should define incident declaration, containment, forensics, legal escalation, communication and notification, with thresholds ensuring that material incidents reach leadership and regulators promptly [8].

3.6 Recover

Recover should use validated backups, tested restoration priorities and alternative procedures, followed by reviews that convert lessons into stronger controls and training [10].

Figure 11. Strategic alignment using the NIST Cybersecurity Framework 2.0.



The programme supports trust, compliance, service continuity and resilience. Investment should be risk-based, with the strongest controls applied to systems containing sensitive identity data.

4. Risk Management and Assessment

Risk management directs resources towards the assets creating the greatest exposure. Equifax shows how an unpatched internet-facing application containing sensitive data becomes a critical business risk when ownership, remediation and verification fail [17].


4.1 Risk assessment method

Risk assessment identifies assets, threats and vulnerabilities. Likelihood considers exploit availability, exposure, attacker capability and existing controls; impact covers financial, regulatory, operational, customer and reputational harm.

$$\text{Risk} = \text{Likelihood} \times \text{Impact}$$

Each risk requires a rating, treatment, named owner and deadline, with reassessment as threats, configurations and priorities change [9].

Figure 12. Cybersecurity risk-assessment process.



4.2 Key Equifax-related risks

Key Equifax-related risks include unpatched external applications, incomplete inventories, unclear ownership, weak segmentation, excessive privilege, monitoring failures, third-party weaknesses, legacy systems, unencrypted data, certificate failures and poor escalation [17].


4.3 Example risk register

Tab.6

Risk	Likelihood	Impact	Priority	Mitigation
Unpatched internet-facing application	High	Critical	Critical	Automated scanning, emergency patching and verification
Weak segmentation	High	High	High	Microsegmentation and Zero Trust access

Risk	Likelihood	Impact	Priority	Mitigation
Excessive privileges	Medium	High	High	Least privilege and Privileged Access Management
Third-party software weakness	Medium	High	High	Supplier assessment and dependency monitoring
Monitoring failure	Medium	Critical	High	SIEM, EDR and alert-health monitoring
Data exfiltration	High	Critical	Critical	Encryption, DLP and behavioural analytics

Figure 13. Illustrative likelihood–impact matrix for Equifax-related risks.



4.4 Risk treatment

Risk may be avoided, reduced, transferred or accepted. Transfer does not remove accountability, and acceptance is appropriate only when residual exposure is within the approved risk appetite [9].

Critical internet-facing vulnerabilities should normally be reduced immediately. Temporary exceptions require CISO approval, compensating controls and an expiry date; priority should reflect exploitability, exposure, data sensitivity and business importance, not severity alone.

5. Compliance, Incident Response and Protection Controls

5.1 Regulatory alignment

A legal-obligations register should cover data-protection law, the Fair Credit Reporting Act, the Gramm–Leach–Bliley Act Safeguards Rule, state notification rules and overseas duties. The Safeguards Rule requires covered institutions to notify the FTC within 30 days of discovering certain events involving at least 500 consumers' unencrypted information [18].

The notification procedure should record the incident, determine affected data and severity, involve Legal and Privacy, meet regulatory deadlines, inform individuals when required and preserve evidence of decisions, findings and actions.

5.2 Incident response plan

The response plan should cover Preparation, Detection and Analysis, Containment, Eradication, Recovery and Lessons Learned [8]. The CISO should command the response, supported by a manager coordinating the SOC, IT, Legal, Privacy, Communications and senior management.

It should define severity, declaration authority, evidence handling, escalation and communication without allowing notification duties to delay containment.




Figure 14. Integrated incident-response and notification process.

5.3 Disaster recovery and business continuity

Business-continuity plans should define recovery-time and recovery-point objectives. Protected backups must be validated through restoration tests, prioritising identity services, customer-facing systems and critical databases, with alternative procedures and exercises supporting continuity [10].

5.4 Technical protection controls

The technical baseline should combine vulnerability scanning, EDR, SIEM, SOAR, intrusion detection and secure configuration. Encryption, MFA, PAM, application testing, microsegmentation and DLP should reduce unauthorised access, software weakness, lateral movement and data loss [11].

5.5 Connection with Zero Trust

Zero Trust evaluates access through identity, resource sensitivity, device posture and contextual risk. Least privilege, segmentation, behavioural monitoring and session revocation should protect critical databases [3].

5.6 Reducing dwell time and organisational impact

SIEM and EDR can expose and isolate compromise earlier; scanning identifies missing patches; SOAR accelerates containment; segmentation limits lateral movement; and encryption, MFA and tested recovery reduce impact [10].

Automation should reflect confidence and business impact: low-risk actions may be automatic, reversible actions monitored, and disruptive actions approved by an analyst.

6. Recommended Cybersecurity Programme Diagram

The programme links Board oversight, cyber-risk governance, operational delivery and continuous improvement. The Board should treat cybersecurity as enterprise risk, approve risk appetite, review major incidents and align investment with business priorities [5].

The CISO and Cyber Risk Governance function should translate that direction into policies, risk treatment, reporting and assurance organised around the six CSF 2.0 functions [5].





Figure 15. Recommended cybersecurity programme aligned with NIST CSF 2.0.

Govern maintains policy and the risk register; Identify covers inventories and impact. Protect covers IAM, MFA and patching. Detect uses SIEM and EDR; Respond coordinates incidents and regulators; Recover provides tested restoration. Reporting should connect the Board Risk Committee, CEO and CISO with Security Architecture, SOC, Vulnerability Management, DevSecOps, Compliance and Incident Response.

Figure 16. Supporting cybersecurity reporting structure.



This structure combines central accountability with specialised operational delivery and a clear escalation route to executive and Board oversight.

7. Performance Measurement

The cybersecurity programme should be evaluated through measurable indicators that demonstrate whether controls are reducing risk and improving resilience. Key measures include the percentage of critical assets inventoried, the percentage of internet-facing systems scanned, the critical patch compliance rate, mean time to detect, mean time to contain, MFA coverage and the number of privileged accounts. Additional indicators should include overdue critical vulnerabilities, the percentage of systems sending logs to SIEM, incident-response exercise results, recovery-time performance and approved policy exceptions [12].

Tab.7

Indicator	Outcome measured
Critical assets inventoried (%)	Asset visibility and ownership completeness
Internet-facing systems scanned (%)	External attack-surface coverage
Critical patch compliance rate	Timely reduction of exploitable exposure
Mean time to detect	Detection speed
Mean time to contain	Containment speed and impact limitation
Accounts protected by MFA (%)	Strength of identity protection
Number of privileged accounts	Scope of high-risk access

Indicator	Outcome measured
Overdue critical vulnerabilities	Unmanaged exposure beyond deadline
Systems sending logs to SIEM (%)	Monitoring coverage and telemetry completeness
Incident-response exercise results	Operational preparedness
Recovery-time performance	Ability to restore within target
Approved policy exceptions	Control deviations and accepted residual risk

Metrics should focus on outcomes, such as critical vulnerabilities remediated on time, verified monitoring coverage, detection accuracy and reduced dwell time. Trends should be reported to the CISO and Board, with major failures escalated immediately.

Overall Conclusion

This project demonstrates that modern cybersecurity can no longer depend primarily on traditional perimeter-based controls. Cloud computing, hybrid work, mobile access and distributed applications weaken the assumption that users and devices inside an organisational network can automatically be trusted. Task 1 established that Zero Trust provides a more suitable model through continuous verification, least-privilege access, identity-based controls, device-posture assessment, microsegmentation, continuous monitoring and automated containment [1].

Zero Trust is not a single product or an immediate replacement for all traditional controls. Its effectiveness depends on organisational maturity, accurate asset inventories, reliable identity governance, integrated technologies, skilled personnel and sustained investment. Traditional firewalls, filtering and secure connectivity remain useful within a multilayered strategy [3].

Task 2 applied these principles to the Equifax breach. The analysis showed that the incident was not caused only by failure to install an Apache Struts patch. It resulted from interconnected weaknesses in governance, accountability, asset visibility, vulnerability ownership, monitoring, segmentation, compliance management and incident escalation [17]. Preventing similar incidents therefore requires a coordinated cybersecurity programme rather than an isolated technical solution.

The proposed programme combines Board oversight, an independent and empowered CISO, clear operational responsibilities and measurable policies. It uses the NIST CSF 2.0 functions to connect cybersecurity activities with organisational objectives, regulatory duties and resilience [5]. Zero Trust principles strengthen this programme by restricting access, assessing contextual risk, monitoring behaviour and limiting lateral movement [1].

The central conclusion is that effective cybersecurity depends on the interaction between governance, people, processes and technology. Controls must be verified, continuously monitored and supported by clear ownership and escalation. Performance should be judged through outcomes such as timely remediation, reduced detection and containment times, complete monitoring coverage and successful recovery testing.

Cybersecurity should therefore be managed as an enterprise-wide, continuously improving risk capability rather than solely as an IT responsibility

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.
- [2] Cybersecurity and Infrastructure Security Agency, “Zero Trust Maturity Model, Version 2.0,” Apr. 2023. [Online]. Available: <https://www.cisa.gov/resources-tools/resources/zero-trust-maturity-model>.
- [3] S. Rose, O. Borchert, A. Kerman, M. Souppaya, et al., “Implementing a Zero Trust Architecture: High-Level Document,” NIST Special Publication 1800-35, Jun. 2025, doi: 10.6028/NIST.SP.1800-35.
- [4] R. Chandramouli and Z. Butcher, “A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Location Environments,” NIST Special Publication 800-207A, Sep. 2023, doi: 10.6028/NIST.SP.800-207A.
- [5] C. Pascoe, S. Quinn, and K. Scarfone, “The NIST Cybersecurity Framework (CSF) 2.0,” NIST CSWP 29, Feb. 2024, doi: 10.6028/NIST.CSWP.29.
- [6] Joint Task Force, “Security and Privacy Controls for Information Systems and Organizations,” NIST Special Publication 800-53, Revision 5, Sep. 2020, updated Nov. 2023, doi: 10.6028/NIST.SP.800-53r5.
- [7] D. Temoshok, J. Fenton, Y.-Y. Choong, N. Lefkowitz, A. Regenscheid, R. Galluzzo, and J. Richer, “Digital Identity Guidelines: Authentication and Authenticator Management,” NIST Special Publication 800-63B-4, Jul. 2025, doi: 10.6028/NIST.SP.800-63b-4.
- [8] A. Nelson, S. Rekhi, M. Souppaya, and K. Scarfone, “Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile,” NIST Special Publication 800-61, Revision 3, Apr. 2025, doi: 10.6028/NIST.SP.800-61r3.
- [9] National Institute of Standards and Technology, “Guide for Conducting Risk Assessments,” NIST Special Publication 800-30, Revision 1, Sep. 2012, doi: 10.6028/NIST.SP.800-30r1.
- [10] M. Swanson, P. Bowen, A. W. Phillips, D. Gallup, and D. Lynes, “Contingency Planning Guide for Federal Information Systems,” NIST Special Publication 800-34, Revision 1, May 2010, doi: 10.6028/NIST.SP.800-34r1.
- [11] M. Souppaya, K. Scarfone, and D. Dodson, “Secure Software Development Framework (SSDF) Version 1.1: Recommendations for Mitigating the Risk of Software Vulnerabilities,” NIST Special Publication 800-218, Feb. 2022, doi: 10.6028/NIST.SP.800-218.
- [12] M. Souppaya and K. Scarfone, “Guide to Enterprise Patch Management Planning: Preventive Maintenance for Technology,” NIST Special Publication 800-40, Revision 4, Apr. 2022, doi: 10.6028/NIST.SP.800-40r4.
- [13] Joint Task Force, “Assessing Security and Privacy Controls in Information Systems and Organizations,” NIST Special Publication 800-53A, Revision 5, Jan. 2022, doi: 10.6028/NIST.SP.800-53Ar5.

- [14] Federal Trade Commission, “Equifax Data Breach Settlement,” 2019. [Online]. Available: <https://www.ftc.gov/enforcement/refunds/equifax-data-breach-settlement>.
- [15] U.S. House of Representatives, Committee on Oversight and Government Reform, “The Equifax Data Breach,” Majority Staff Report, Dec. 2018. [Online]. Available: <https://oversight.house.gov/wp-content/uploads/2018/12/Equifax-Report.pdf>.
- [16] U.S. Government Accountability Office, “Data Protection: Actions Taken by Equifax and Federal Agencies in Response to the 2017 Breach,” GAO-18-559, Aug. 2018. [Online]. Available: <https://www.gao.gov/products/gao-18-559>.
- [17] Federal Trade Commission, “Equifax, Inc.: Complaint for Permanent Injunction and Other Equitable Relief,” Jul. 2019. [Online]. Available: https://www.ftc.gov/system/files/documents/cases/172_3203_equifax_complaint_7-22-19.pdf.
- [18] Federal Trade Commission, “FTC Safeguards Rule: What Your Business Needs to Know,” and 16 C.F.R. Part 314, Standards for Safeguarding Customer Information. [Online]. Available: <https://www.ftc.gov/business-guidance/resources/ftc-safeguards-rule-what-your-business-needs-know>.