

8/10/2026

Data Security and Compliance

Assessment 1: Research and case scenario in Data Security and Compliance

Module Name: Data Security and Compliance

Module Code: ACCA 5034

Student Name: Evgeni Veselinov Iliev

Student ID: 2424555

Lecturer's Name: Dr. Ayman El Hajjar

Table of Contents

Task 1. The Zero Trust Architecture.....	2
1.1. Introduction	2
1.2. Principles and operation of Zero Trust	3
1.3. Comparison with the traditional perimeter model	4
1.4. Creating an environment that anticipates threats	5
1.5. Protecting users, devices, applications and data.....	6
1.6. Detection and automated response	7
1.7. Implementation challenges and limitations	8
1.8. Practical migration and measurement approach	9
1.9. Critical evaluation and conclusion	10
Task 2. Equifax Case Scenario	11
2.1. Case context and programme objective	11
2.2. Organisational structure and accountability	11
2.3. Policies and procedures	13
2.4. Governance and strategic alignment.....	14
2.5. Risk management and assessment	15
2.6. Compliance, incident response and protection controls.....	16
2.6.1 Regulatory alignment and notification	16
2.6.2 Incident response and disaster recovery	16
2.6.3 Technical safeguards and reduction of dwell time	17
2.7. Assurance and continuous improvement.....	18
2.8. Conclusion.....	18
References	19

Task 1. The Zero Trust Architecture

1.1. Introduction

The traditional approach to organisational security was built around a trusted internal network protected by firewalls, virtual private networks and other perimeter controls. This made sense when most staff, devices and applications were located in one building. It is less suitable for organisations that now depend on cloud services, home working, mobile devices and third-party access. Once an attacker passes the perimeter, a traditional network can also give them more freedom than they should have. Zero Trust Architecture (ZTA) responds to this weakness by removing automatic trust based on location and making access decisions about the user, device, resource and current level of risk [1].

This paper argues that ZTA can improve prevention, detection and response, particularly by limiting lateral movement and continuously checking access. However, it is not a product that can simply be installed. Its success depends on accurate asset information, mature identity management, usable policies, reliable monitoring and strong governance. Therefore, ZTA is best treated as a long-term operating model that supports the whole security lifecycle, rather than a replacement name for existing access-control tools.

1.2. Principles and operation of Zero Trust

NIST describes zero trust as a set of principles rather than one fixed design [1]. The central principle is that no account or device receives implicit trust because it is inside a corporate network or owned by the organisation. All data sources and services are treated as resources, all communications are secured, and access is granted to a specific resource for a specific session. Decisions should use dynamic information such as identity, role, device health, time, location, application sensitivity and unusual behaviour. Least privilege is important because successful authentication should not give a person access to everything connected to the network.

The logical design has three main functions. A Policy Engine makes the access decision, a Policy Administrator establishes or ends the session, and a Policy Enforcement Point carries out the decision between the requester and the resource [2]. These functions rely on information from identity services, endpoint tools, asset inventories, threat intelligence and activity logs. For example, an employee may normally access a payroll application from a managed laptop. If the same account suddenly connects from an unmanaged device with missing security updates, the policy could require stronger authentication, allow read-only access or reject the request. Trust is therefore temporary and based on evidence.

The phrase “never trust, always verify” is useful but can be misleading if taken literally. Organisations must still trust identity providers, device records, administrators and policy systems. ZTA does not remove trust; it reduces, checks and distributes it more carefully. This creates better control, but it also makes the accuracy and resilience of these supporting systems critical.

1.3. Comparison with the traditional perimeter model

Table 1 compares the main differences. A perimeter is still useful for blocking unwanted traffic, so zero trust should not mean removing firewalls. The real change is that network location is no longer enough to prove that an access request is safe.

Table 1. Traditional perimeter security compared with Zero Trust Architecture

Area	Traditional model	Zero trust model	Critical effect
Trust	Internal network is treated as safer	No automatic trust from location	Reduces the value of an initial foothold
Access	Broad network or VPN access	Per-session access brokered to individual resources	Limits unnecessary exposure for local and remote users
Decision data	Mainly credentials and network location	Identity, device posture, resource and behaviour	Gives richer decisions but needs reliable data
Monitoring	Often focused on network boundaries	Continuous user, endpoint and resource telemetry	Improves detection but increases cost and privacy concerns
Failure risk	Perimeter bypass may expose many systems	Policy or identity failure may affect many decisions	New control points require strong resilience

Google's BeyondCorp programme shows that this idea can work at scale. It moved access decisions away from the corporate network and used user credentials, device information and an access proxy to protect individual services [3]. This is valuable practical evidence, but it should not be copied without thought. Google has engineering resources and internally developed systems that many smaller organisations do not have. The case demonstrates feasibility, not that every organisation will gain the same result or follow the same route.

1.4. Creating an environment that anticipates threats

Anticipation begins with governance and knowledge of the environment. Leaders should identify important business services, data owners, legal duties, threat scenarios and acceptable levels of risk. NIST Cybersecurity Framework (CSF) 2.0 places Govern at the centre of its six functions because roles, policy and risk appetite shape every other activity [4]. A ZTA programme should therefore have executive sponsorship, named owners and measurable outcomes, such as the percentage of privileged accounts using phishing-resistant multi-factor authentication (MFA), the number of unmanaged assets, and the time taken to remove access after a role change.

Accurate inventories are essential. An organisation cannot create a policy for a resource it does not know exists. Hardware, software, service accounts, cloud workloads, certificates, data sets and dependencies should be continuously discovered and linked to business owners. Data should also be classified so that stronger controls are applied to personal, financial or commercially sensitive information. This avoids spending equal effort on low-risk and high-risk services.

Threat intelligence and vulnerability information can then change policy before an incident occurs. If active exploitation is reported for software used by an organisation, vulnerable assets can be found, prioritised and temporarily isolated until patched. Tabletop exercises and attack simulation can test whether policies behave as expected. The Cybersecurity and Infrastructure Security Agency (CISA) maturity model is helpful because it divides progress across identity, devices, networks, applications and workloads, and data, with visibility, analytics, automation and governance cutting across them [5]. This discourages a narrow project that modernises login controls while leaving data and legacy servers exposed.

1.5. Protecting users, devices, applications and data

Protection should start with identity. Each user and service account needs a unique identity, MFA and access based on current job requirements. Privileged administrators should use separate accounts with time-limited elevation. Joiner, mover and leaver processes must quickly create, change and remove permissions. These measures reduce the chance that an old or stolen account will provide lasting access.

Device posture adds another layer. Managed laptops can be checked for encryption, endpoint protection, supported software and current patches before access is allowed. An unmanaged personal device might receive a restricted web session rather than a direct connection. This is stronger than a virtual private network (VPN), which can authenticate a user and then place their device onto a wider network. However, device checks must be proportionate. Blocking every slightly out-of-date device could stop important work and encourage staff to find workarounds.

Micro segmentation and application-level gateways restrict communication between systems. If one web server is compromised, it should only reach the services needed for its purpose, not every database. Sensitive data should be encrypted in transit and at rest, with keys separated from the encrypted information. Data Loss Prevention (DLP) can restrict copying or large transfers, while query limits can identify activity that is technically authorised but abnormal. These controls are important because ZTA mainly governs access; it does not automatically correct insecure code, excessive data retention or weak encryption.

NIST's 2025 implementation guide used 24 collaborators to build 19 example ZTA solutions, showing that several approaches can combine identity governance, micro segmentation, software-defined perimeter and Secure Access Service Edge (SASE) technologies [6]. This supports a flexible, risk-based design. It also shows why organisations should define requirements before selecting suppliers, as no single product provides the whole architecture.

1.6. Detection and automated response

Continuous monitoring turns zero trust from a one-time login check into an adaptive control. Identity logs, endpoint detection and response (EDR), cloud events, network metadata and application activity should feed a Security Information and Event Management (SIEM) platform. Analytics can compare activity with expected behaviour. Examples include an account downloading an unusual amount of data, a server contacting an unknown destination, or a privileged login from an unhealthy device.

The Policy Engine can use this information immediately. Low-risk events may only increase logging. Medium-risk events could require reauthentication or reduce access. High-confidence events can end a session, disable a token, isolate an endpoint or block a particular connection. Security Orchestration, Automation and Response (SOAR) playbooks can also open a case, preserve evidence and alert the incident team. NIST found that SIEM and SOAR information is most valuable when it reaches policy decision points in real time [6]. This closes the gap between detecting a threat and applying a control.

Automation must still be controlled. A false positive that disables a hospital system, payment service or administrator account may cause serious harm. Playbooks should therefore have approved thresholds, tested rollback steps and human approval for high-impact actions. Responses should be logged so they can be audited and improved. NIST incident guidance also treats Govern, Identify and Protect as preparation for Detect, Respond and Recover, showing that automation cannot compensate for weak planning [7].

1.7. Implementation challenges and limitations

The first challenge is poor information quality. Asset registers are often incomplete, identities are duplicated and application owners are unclear. A wrong device record can block a genuine employee or approve a compromised machine. NIST states that an organisation needs a solid continuous-monitoring capability before implementing ZTA [2]. This means organisations with weak basic controls may need to improve them before advanced dynamic policies provide value.

Legacy systems create a second challenge. Older applications may not support modern authentication, short-lived tokens or fine-grained authorisation. Replacing them can be costly and disruptive. A gateway can protect some legacy services, but it may only provide coarse access and become another point of failure. Migration will normally be gradual, leaving a hybrid environment in which attackers target the parts that still depend on traditional trust [1].

Integration is also difficult. NIST's practical builds found that products from different vendors did not always share endpoint compliance or risk information, and multiple policy decision points could hold inconsistent views [6]. Proprietary interfaces can cause vendor lock-in and make troubleshooting harder. Central identity and policy services may also become high-value targets. If they are compromised, misconfigured or unavailable, access decisions across the organisation can be affected. Resilient design therefore needs replicated services, protected administration, change approval, audit logs and tested continuity arrangements.

People and privacy cannot be ignored. Frequent challenges may create security fatigue, while detailed monitoring can feel intrusive, particularly on personal devices. Data collection should be limited to what is needed for a stated security purpose, retained for a defined period and protected from misuse. Training must explain why access has changed and provide a quick route for correcting mistakes. Without this, users may resist the programme even if its technical design is sound.

Finally, the evidence base has limits. NIST notes that common measures for comparing ZTA with older strategies still need development [1]. Reduced lateral movement, fewer standing privileges, faster containment and lower exposure are sensible measures, but organisations should not claim that the architecture prevents all breaches. Stolen authorised sessions, vulnerable applications, malicious administrators and mistakes remain possible.

1.8. Practical migration and measurement approach

A sensible migration should begin with one important but manageable business service. The organisation should map its users, devices, data flows and dependencies, then record the current controls and risks. It can create a target access policy, introduce stronger authentication and place a Policy Enforcement Point in front of the service. Running the new policy in monitoring mode first allows the team to see which legitimate activities would be blocked. This evidence can be used to correct identities, device records and access rules before enforcement. The approach is slower than a large technology rollout, but it reduces the chance of disrupting the business.

Progress should be measured through security and operational outcomes. Useful measures include the percentage of high-value applications behind per-session access, privileged permissions removed, unmanaged devices attempting access, time taken to revoke a risky session and success rate of recovery tests. Helpdesk calls and failed legitimate access should also be measured. A programme that blocks attacks but regularly prevents employees from working is not mature. Likewise, a fall in alerts may mean that risk has reduced, but it could also mean that telemetry has failed. Measures need explanation and independent checks rather than being accepted at face value.

The target state should not be complete removal of the perimeter. Firewalls, secure configuration, patch management, backups and incident response still provide important layers. The aim is to stop one layer from carrying too much trust. Each migration stage should therefore include a threat simulation and a failure test. The team should confirm that a compromised account cannot reach unrelated resources, while also proving that critical access continues if a policy service becomes unavailable. Lessons should update both policy and business continuity plans.

This staged method also supports financial control. Leaders can compare the cost of each phase with reduced standing access, faster containment and retired legacy connections. If a proposed tool does not improve an agreed outcome, the organisation should question the purchase rather than treating it as progress. This keeps zero trust connected to risk and business value.

1.9. Critical evaluation and conclusion

ZTA is more suitable than a perimeter-only model for modern organisations because users, data and systems no longer sit within one dependable boundary. Its strongest benefit is not perfect prevention but reduced impact. A compromised account or server should face repeated checks and narrow permissions, making movement and data theft harder. Continuous telemetry can also shorten detection and enable a response during an attack rather than after it.

However, its effectiveness depends on implementation quality. The architecture may increase complexity, reliance on identity services and the consequences of incorrect automated decisions. A mature firewall-based environment with good segmentation, patching and monitoring could be safer than a badly designed project labelled “zero trust”. The most reasonable approach is an incremental migration focused on high-value data and risky workflows. Organisations should establish governance, map assets and data, strengthen identity, introduce least privilege and monitoring, then automate low-risk responses before progressing to disruptive actions. ZTA should therefore be judged by measurable risk reduction and business resilience, not by the number of new products purchased.

Compliance can support this work, although it should not become the only reason for adoption. Evidence from access reviews, policy decisions, device checks and response logs can demonstrate that controls are operating. At the same time, passing an audit only proves that selected requirements were met at a point in time. Attackers adapt and business systems change. The organisation still needs continuous risk assessment, technical testing and honest reporting of control failures. In this sense, zero trust is most effective when it changes daily decisions as well as formal documentation. It creates a habit of asking what evidence justifies access and what damage is possible if that decision is wrong.

Task 2. Equifax Case Scenario

2.1. Case context and programme objective

The breach exposed personal data belonging to approximately 147 million people, but it was not caused by one missed patch alone. Attackers exploited CVE-2017-5638 in an internet-facing Apache Struts application, a flaw that allowed remote command execution [8]. The United States Government Accountability Office (GAO) found weaknesses in identification, detection, segmentation and data governance [9]. Attackers used unencrypted credentials to move beyond the dispute portal, reached 48 additional databases, ran about 9,000 queries and removed data over encrypted connections for approximately 76 days [9]. The United States Senate investigation also highlighted an incomplete asset inventory, ineffective scanning and an expired monitoring certificate [10].

The proposed programme has two objectives: prevent a similar initial compromise and limit the damage if prevention fails. It combines NIST CSF 2.0 with the Zero Trust principles assessed in Task 1. This is appropriate because CSF 2.0 provides governance and lifecycle outcomes, while ZTA gives a practical rule for access: no user, device or workload receives broad trust merely because it has entered the network.

2.2. Organisational structure and accountability

Cybersecurity should be treated as an enterprise risk, not a technical issue owned only by the IT department. The board Risk and Audit Committee should approve risk appetite, challenge overdue high risks and receive a quarterly dashboard. The Chief Information Security Officer (CISO) should report to the Chief Executive Officer (CEO), with independent access to the committee. This gives security a voice separate from the Chief Information Officer (CIO), whose delivery and availability targets can conflict with urgent patching. GAO reported that Equifax introduced direct CISO-to-CEO reporting after the breach, which supports this allocation [9].

Table 2. Governance roles, reporting lines and accountability

Role	Reporting or access	Main accountability
Board Risk and Audit Committee	Receives reports from the CISO, DPO and Internal Audit	Approves risk appetite; challenges overdue critical risks; reviews quarterly trends and control evidence
CEO	Accountable to the board	Receives direct CISO and CIO reports; resolves conflicts between security, availability and delivery
CISO	Reports to the CEO; direct access to the board committee	Owns security strategy, policy and risk reporting; escalates risks that have not been accepted
CIO and service or data owners	Operational reporting lines	Operate secure services, meet patch targets, provide evidence and formally accept residual risk
Security Architecture, Product Security, Vulnerability Management and SOC	Under CISO and security leadership	Design segmentation and SSDLC controls; maintain asset visibility, monitoring and incident response
DPO and Internal Audit	Independent access to senior management and the board	The DPO advises on privacy and notification; Internal Audit tests control evidence

The CISO owns security strategy, policy and risk reporting but should not personally own every control. The CIO remains accountable for operating secure infrastructure and meeting patch targets. A Head of Security Architecture designs segmentation and Zero Trust policies. A Product Security Lead oversees the secure software development lifecycle (SSDLC), including dependency checks and release gates. A Vulnerability and Asset Manager maintains discovery, scanning and remediation assurance. The Security Operations Centre (SOC) monitors events and leads containment. A Data Protection Officer (DPO) advises independently on privacy, notification and data minimisation. Business service and data owners accept residual risk because they understand operational impact. Internal Audit reports to the board committee and tests whether evidence supports management's claims.

This allocation follows a three-lines approach: operational teams own and operate controls; security and risk teams set standards and challenge performance; Internal Audit provides independent assurance. Any exception that leaves a critical internet-facing vulnerability open beyond its deadline must be accepted by the relevant executive and CISO, not silently carried by an engineer.

2.3. Policies and procedures

Policies must state an owner, a measurable deadline, required evidence and an escalation route. Equifax reportedly ordered vulnerable systems to be patched within 48 hours but did not verify completion [11]. The redesigned process therefore separates “patch deployed” from “remediation verified”. NIST also defines enterprise patch management as identifying, prioritising, acquiring, installing and verifying patches [12]. Table 2 turns this principle into auditable rules.

Table 2. Core security policies and assurance measures

Policy and owner	Measurable requirement	Evidence and escalation	Equifax weakness addressed
Asset visibility — CIO	Continuous discovery; every production asset, application, certificate and owner recorded within 24 hours	Daily reconciliation; unknown internet asset isolated; monthly completeness report	Unknown Struts system and expired certificate
Vulnerability management — Vulnerability Manager	Authenticated internal scans weekly; external attack-surface scan daily; active exploitation and critical internet flaws triaged within 4 hours	Scanner coverage, false-negative review and ticket linkage; gaps escalated to CISO	Scan failed to identify vulnerable portal
Patch management — service owner	Critical exploited/internet-facing flaws mitigated within 24 hours and patched within 48 hours; high risk within 7 days; medium within 30 days	Independent rescan and configuration proof; overdue items go to executive risk register	Order issued but completion not checked
SSDLC — Product Security Lead	Threat model for high-risk changes; peer review, SAST, DAST, dependency and secret scanning in the pipeline; no critical finding at release	Signed build record; exceptions expire within 30 days; quarterly sample by Internal Audit	Vulnerable component and weak development oversight
Monitoring — SOC Manager	EDR and central logging on at least 98% of supported assets; security certificates monitored 60, 30 and 7 days before expiry	Daily control-health alerts; missing telemetry treated as an incident	Inspection failed for about ten months
Data protection — DPO and data owner	Encrypt sensitive data in transit and at rest; no plaintext credentials; quarterly access review; retention based on purpose	Key and access logs, DLP events and deletion evidence; material gaps to Risk Committee	Unencrypted credentials and excessive reachable data

The timings are intentionally risk based. A fixed 30-day target for every issue would leave an easily exploited public system open too long, while emergency patching of every low-risk flaw would cause unnecessary outages. The CISA Known Exploited Vulnerabilities catalogue should be an input to priority because it identifies flaws with evidence of active exploitation [14]. Where a patch cannot be applied, the owner must use a time-limited compensating control such as disabling the feature, a web application firewall rule or network isolation. The exception must record business justification, test results and an expiry date.

2.4. Governance and strategic alignment

NIST CSF 2.0 is suitable because its outcomes can be adapted rather than followed as a rigid checklist [4]. The organisation should produce a Current Profile, a Target Profile and a costed improvement plan. Cyber risks then appear in the enterprise risk register alongside financial, legal and operational risks. This connects security spending to consumer trust, service availability and regulatory exposure.

Table 3. Strategic alignment using NIST CSF 2.0

CSF function	Programme outcome	Business measure
Govern	Board-approved appetite, ownership, policy and supplier oversight	No unowned critical risk; quarterly board challenge recorded
Identify	Live asset, data, dependency and risk knowledge	Inventory coverage and critical-service dependency maps
Protect	Least privilege, secure development, patching and encryption	Patch compliance, MFA coverage and access-review completion
Detect	Healthy telemetry and tested analytics	Mean time to detect; percentage of assets reporting to SIEM/EDR
Respond	Practised decisions, containment and communication	Mean time to contain; exercise actions closed on time
Recover	Restorable services and trusted communications	Recovery time/objective tests and successful restore rate

The board dashboard should show trends and exceptions, not only a single “green” score. A quarterly cycle should compare the Current and Target Profiles, approve priorities and track benefits. Business Continuity Planning (BCP) is part of this strategy. Critical credit-reporting services need defined Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs), tested alternatives and clear manual procedures. Trust is also an outcome: accurate, timely communication and evidence of control testing protect customers and partners more effectively than broad claims that the organisation is secure.

2.5. Risk management and assessment

Risks should be scored using likelihood and impact on a five-point scale, producing an inherent rating before controls and a residual rating after treatment. Ratings guide attention but do not replace judgement. A lower-frequency event can remain a priority when it exposes permanent identity information or a critical public service. Table 4 focuses on risks shown by Equifax and on modern supply-chain exposure.

Table 4. High-level cybersecurity risk register

Risk and case evidence	Inherent rating	Priority treatment	Residual rationale
Exploitation of an internet-facing application; Struts allowed remote commands [8], [9]	Likelihood 5, impact 5 — Critical	Daily discovery, WAF, secure build pipeline, 24-hour mitigation, 48-hour verified patch	Medium: exposure remains, but attack window and vulnerable releases are reduced
Asset or certificate missing from governance; portal and inspection certificate were not controlled [9], [10]	5 × 5 — Critical	Automated discovery, named owners, configuration database reconciliation and expiry automation	Low–medium: unknown assets trigger isolation and management alert
Lateral movement through weak segmentation and stored credentials [9]	4 × 5 — Critical	Application-level access, network micro segmentation, secrets vault and least privilege	Medium: a compromised service has fewer reachable targets
Undetected exfiltration; encrypted traffic and 9,000 queries blended with normal use [9]	4 × 5 — Critical	EDR, SIEM analytics, database activity monitoring, DLP, rate limits and egress controls	Medium: small theft may occur, but abnormal volume and behaviour should be contained sooner
Third-party or open-source component compromise	4 × 4 — High	Supplier due diligence, software bill of materials, dependency scanning, signed builds and contract notification clauses	Medium: provenance improves, but supplier assurance cannot remove all inherited risk
Zero Trust control-plane failure or false-positive automation	3 × 5 — High	Replicated policy services, break-glass process, change approval, tested rollback and human approval for disruptive actions	Low–medium: resilience prevents one decision service from stopping critical operations

The first four risks receive immediate funding because they combine high likelihood with severe privacy impact and match the known attack path. Supply-chain risk is next because modern applications rely on external code and services. Control-plane risk is included to avoid creating a new single point of failure while correcting the old perimeter model. Risk owners must review critical items monthly, test control effectiveness, and record residual risk rather than assuming that deployment equals protection.

2.6. Compliance, incident response and protection controls

2.6.1 Regulatory alignment and notification

Legal and compliance teams should maintain a jurisdiction matrix showing affected data, regulator, notification threshold, deadline and responsible decision maker. For European Union personal data, GDPR Article 33 normally requires notification to the supervisory authority within 72 hours of awareness where the breach creates risk to individuals [15]. A US-listed company must also consider the Securities and Exchange Commission requirement to file within four business days after determining that an incident is material [16]. Consumer reporting activities bring additional US duties and regulatory oversight. The Equifax settlement reached by the Federal Trade Commission, Consumer Financial Protection Bureau and states was worth at least \$575 million and required an information-security programme [11].

The incident commander should notify the CISO, DPO, Legal Counsel and materiality committee as soon as personal data or a major service may be affected. Legal Counsel owns the notification timetable but cannot delay technical containment. Decisions, known facts and uncertainty should be recorded. Affected people should receive clear information about the data, likely consequences and protective steps. This avoids both late silence and inaccurate early claims.

2.6.2 Incident response and disaster recovery

The response plan should follow Detect, Respond and Recover, supported by preparation across Govern, Identify and Protect [7]. The SOC validates an alert and sets severity. The incident commander then coordinates containment, evidence preservation, legal assessment, communications and service recovery. For a repeat of Equifax, approved automated actions would block the exploit pattern, isolate the web workload, revoke exposed secrets and restrict database routes. A human would approve taking a public service fully offline unless immediate harm made pre-authorised shutdown necessary.

Playbooks should cover vulnerable public applications, stolen privileged accounts, large data exports, ransomware and supplier compromise. Exercises should run twice each year, with one technical simulation and one board-level scenario. Lessons must become assigned actions with deadlines. Disaster recovery should use encrypted, immutable backups separated from normal administrator accounts. Quarterly restoration tests must prove that critical services meet their RTO and RPO; a backup that has never been restored is only an assumption.

2.6.3 Technical safeguards and reduction of dwell time

Table 5 shows how the main safeguards work together. The aim is defence in depth: if patching fails, segmentation limits reach; if access is abused, monitoring detects behaviour; if detection is confident, automation contains it.

Table 5. Protection, detection and response controls

Safeguard	Preventive or detective value	Automated response and impact
Continuous asset and vulnerability scanning	Finds public assets, outdated components and configuration drift	Opens an owned ticket; isolates an unknown critical asset; verifies remediation by rescan
EDR and SIEM	Detects command execution, unusual processes, identity events and missing telemetry	Isolates endpoint, revokes session and preserves evidence, reducing time from alert to containment
Database monitoring and DLP	Detects abnormal queries, sensitive records and unusual transfer volume	Rate-limits query, blocks export and alerts data owner, limiting records exposed
Microsegmentation and ZTNA	Allows only required service-to-service and user-to-application paths	Removes the affected route without closing the whole network, reducing lateral movement
Encryption and secrets management	Protects stored data and removes plaintext database credentials	Rotates exposed secrets and keys; stolen files have less immediate value
Phishing-resistant MFA	Makes password theft insufficient for sensitive access [17]	Step-up authentication or token revocation when device or behaviour risk changes

Encryption keys must be stored and governed separately from encrypted data, following recognised key-management practice [18]. Monitoring coverage also needs a “dead sensor” alert; the expired Equifax certificate shows that installing a tool is not the same as receiving protection. Useful operational targets are: high-severity alert triage within 15 minutes, containment within 60 minutes, at least 98% supported-endpoint telemetry, and monthly testing of one automated playbook. These targets should be adjusted after exercises, but missed targets require analysis rather than being quietly reset.

2.7. Assurance and continuous improvement

The programme needs an assurance cycle so that targets do not become paperwork. Control owners should check their evidence each month, while the second-line security risk team challenges overdue actions and samples important results. Internal Audit should independently test the highest risks each year, including whether assets can be discovered, patches are genuinely installed, access reviews remove permissions and backups can be restored. Penetration testing should then test complete attack paths rather than only listing technical weaknesses.

Management should use a small set of linked measures. Leading indicators include asset coverage, patch age, privileged MFA, certificate expiry and logging health. Lagging indicators include incidents, dwell time, records exposed and service downtime. Neither group is sufficient alone. Good patch statistics do not prove that scanners can see every asset, while a year without a major incident does not prove that controls are effective. Exceptions and failed tests should therefore appear beside headline percentages in the board report.

Every serious incident, exercise and near miss should lead to a short lessons review. The review must identify a control or decision that needs to change, name an owner and set a date. The CISO should verify closure, and the Risk Committee should challenge repeat findings. This creates a feedback loop between operations and governance. It also supports the principle from Task 1 that access decisions improve as the organisation gathers better evidence about assets, threats and behaviour.

2.8. Conclusion

The proposed programme addresses the Equifax breach as a failure of governance and assurance as well as technology. Direct board visibility, clear control owners, verified patching, live asset knowledge, secure development and risk-based segmentation reduce the chance of entry and movement. Continuous monitoring, tested playbooks and carefully limited automation reduce dwell time when prevention fails. Most importantly, the programme requires evidence that controls are working. This turns cybersecurity from a collection of tools into an accountable system that supports privacy, business continuity and long-term trust.

References

- [1] S. Rose, O. Borchert, S. Mitchell and S. Connelly, *Zero Trust Architecture*, NIST Special Publication 800-207. Gaithersburg, MD, USA: National Institute of Standards and Technology, Aug. 2020, doi: 10.6028/NIST.SP.800-207.
- [2] S. Rose, *Planning for a Zero Trust Architecture: A Planning Guide for Federal Administrators*, NIST Cybersecurity White Paper 20. Gaithersburg, MD, USA: National Institute of Standards and Technology, May 2022, doi: 10.6028/NIST.CSWP.20.
- [3] R. Ward and B. Beyer, "BeyondCorp: A New Approach to Enterprise Security," *login*, vol. 39, no. 6, pp. 6–11, Dec. 2014. [Online]. Available: https://www.usenix.org/system/files/login/articles/login_dec14_02_ward.pdf. [Accessed: 11 Aug. 2026].
- [4] National Institute of Standards and Technology, *The NIST Cybersecurity Framework (CSF) 2.0*, NIST Cybersecurity White Paper 29. Gaithersburg, MD, USA: NIST, Feb. 2024, doi: 10.6028/NIST.CSWP.29.
- [5] Cybersecurity and Infrastructure Security Agency, *Zero Trust Maturity Model, Version 2.0*. Washington, DC, USA: CISA, Apr. 2023. [Online]. Available: <https://www.cisa.gov/resources-tools/resources/zero-trust-maturity-model>. [Accessed: 11 Aug. 2026].
- [6] O. Borchert et al., *Implementing a Zero Trust Architecture: High-Level Document*, NIST Special Publication 1800-35. Gaithersburg, MD, USA: National Institute of Standards and Technology, June 2025, doi: 10.6028/NIST.SP.1800-35.
- [7] 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 Rev. 3. Gaithersburg, MD, USA: NIST, Apr. 2025, doi: 10.6028/NIST.SP.800-61r3.
- [8] National Institute of Standards and Technology, "CVE-2017-5638 Detail," *National Vulnerability Database*. [Online]. Available: <https://nvd.nist.gov/vuln/detail/CVE-2017-5638>. [Accessed: 11 Aug. 2026].
- [9] U.S. Government Accountability Office, *Data Protection: Actions Taken by Equifax and Federal Agencies in Response to the 2017 Breach*, GAO-18-559. Washington, DC, USA: GAO, Aug. 2018. [Online]. Available: <https://www.gao.gov/assets/gao-18-559.pdf>. [Accessed: 11 Aug. 2026].
- [10] U.S. Senate Permanent Subcommittee on Investigations, *How Equifax Neglected Cybersecurity and Suffered a Devastating Data Breach*. Washington, DC, USA, Mar. 2019. [Online]. Available: <https://www.hsgac.senate.gov/subcommittees/investigations/rep/carper-portman-unveil-bipartisan-report-detailing-equifaxes-repeated-failures-to-protect-sensitive-information-of-more-than-145-million-americans/>. [Accessed: 11 Aug. 2026].
- [11] Federal Trade Commission, "Equifax to Pay \$575 Million as Part of Settlement with FTC, CFPB, and States Related to 2017 Data Breach," 22 July 2019. [Online]. Available: <https://www.ftc.gov/news-events/news/press-releases/2019/07/equifax-pay-575-million-part-settlement-ftc-cfpb-states-related-2017-data-breach>. [Accessed: 11 Aug. 2026].
- [12] M. Souppaya and K. Scarfone, *Guide to Enterprise Patch Management Planning: Preventive Maintenance for Technology*, NIST Special Publication 800-40 Rev. 4. Gaithersburg, MD, USA: NIST, Apr. 2022, doi: 10.6028/NIST.SP.800-40r4.
- [13] 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. Gaithersburg, MD, USA: NIST, Feb. 2022, doi: 10.6028/NIST.SP.800-218.
- [14] Cybersecurity and Infrastructure Security Agency, "Known Exploited Vulnerabilities Catalog." [Online]. Available: <https://www.cisa.gov/known-exploited-vulnerabilities-catalog>. [Accessed: 11 Aug. 2026].
- [15] European Parliament and Council of the European Union, "Regulation (EU) 2016/679 (General Data Protection Regulation), Article 33," *Official Journal of the European Union*, 27 Apr. 2016. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>. [Accessed: 11 Aug. 2026].
- [16] U.S. Securities and Exchange Commission, "Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure," Release Nos. 33-11216 and 34-97989, 26 July 2023. [Online]. Available: <https://www.sec.gov/rules-regulations/2023/07/s7-09-22>. [Accessed: 11 Aug. 2026].
- [17] D. Temoshok et al., *Digital Identity Guidelines: Authentication and Authenticator Management*, NIST Special Publication 800-63B-4. Gaithersburg, MD, USA: NIST, July 2025, doi: 10.6028/NIST.SP.800-63B-4.
- [18] E. Barker, *Recommendation for Key Management: Part 1 — General*, NIST Special Publication 800-57 Part 1 Rev. 5. Gaithersburg, MD, USA: NIST, May 2020, doi: 10.6028/NIST.SP.800-57pt1r5.