Essentials for enterprise IT

Government and industry consultation paper

Published: 15 June 2026

Table of contents

Executive Summary **1**

Part 1: Purpose **3**

How to read this document 4

The evolving cyber threat environment 5

What is enterprise IT 5

What has changed 5

Part 2: Framework structure **9**

Cyber security principles 10

Implementation approaches 14

Protect surfaces 14

Implementation levels 15

Understanding the protect surfaces 18

Part 3: Implementing the framework **25**

Options for implementation focus 25

Getting started with implementation 28

Measuring implementation effectiveness 33

Recognising existing Essential Eight investment 42

How the Essentials informs risk-based decisions 44

Glossary **47**

Annex A: Cyber security principles and implementation approaches **50**

Summary 50

Executive Summary

The Essential Eight mitigation strategies is part of the Australian Signals Directorate's (ASD) flagship guidance to help Australian organisations protect themselves from cyber threats. Originally released in 2017, and regularly updated, the Essential Eight was designed to provide guidance for enterprise IT networks. ASD's Essential Eight remains effective in preventing cyber security incidents but the threat environment and technologies have evolved. With this in mind, ASD is shifting to a broader, principles-based 'Essentials' series to respond to contemporary tradecraft and modern enterprise IT environments, including cloud and operational technology (OT). This is to ensure the guidance remains effective, achievable and fit for purpose.

The proposed changes are informed by ASD's unique intelligence insights combined with cyber security incident data, partner feedback and public cyber threat reporting. ASD has taken a principled approach to offer prioritised, threat-informed mitigations for modern IT environments, supported by clear progression pathways and improved usability. ASD recommends that organisations continually review their risk and security settings, to protect themselves from the dynamic cyber security landscape.

The consultation paper presents the proposed principles-based approach to cyber security, while offering flexibility in implementation and a clear and measurable path towards an improved security posture. There are three parts to the paper; 1) Purpose, 2) Framework Structure and 3) Implementing the framework. The proposed framework has been designed around five structural shifts:

- Prescriptive controls to flexible principles
- Fixed maturity to targeted implementation
- Uniform progression to sequenced implementation
- Maintaining a broad threat approach, with increased focus on the level of sophistication of cyber threats.
- Tailored for Commonwealth entities and CI operators but remains relevant for all organisations.

Overall, the consultation paper provides detailed information on the proposed framework to assist stakeholders to understand the need for change, outline how the frameworks is designed and guidance on the implementation. Stakeholders are invited to provide a submission on the Essentials for enterprise IT via ASD's Partner Portal by 12 July.

The Framework structure

The proposed framework sets out a layered approach that provides organisations with a clear picture of what to secure, what good looks like and how to get there, with several case study examples. The framework revolves around defined cyber security principles, followed by implementation approaches and grouped by protect surfaces, as shown in figure 1.

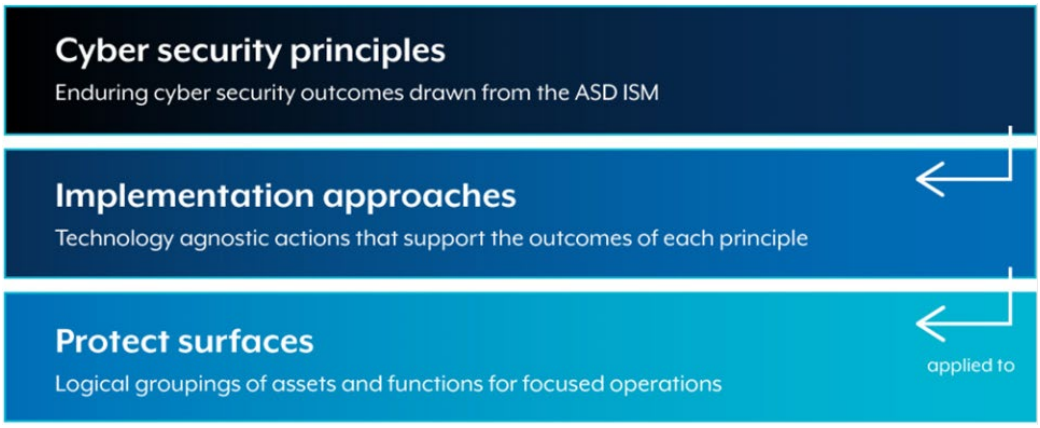


Figure 1 – Framework structure.

Each cyber security principle is drawn from ASD's Information Security Manual (ISM) which is tailored through implementation approaches that prioritise the administrative or technical outcomes required, which are scalable and technology-agnostic. The protect surfaces are a new feature which groups enterprise functions and associated assets that share a common risk profile.

The implementation and measurement approach

The consultation paper offers two options for implementation; organisations can implement a protect surface (option 1) or a principle approach (option 2). Option 1 involves selecting the most appropriate protect surface according to the organisation's risks, and implementing the associated principles. Option 2 adopts the most relevant principles and applies them across all the associated protect surfaces. Under both options, organisations will need to consider their risk and operating environment, and select the appropriate controls from a range of cyber security frameworks, including but not limited to the ISM. The consultation paper outlines the steps involved in this process.

A key change is the removal of maturity levels, replaced with implementation levels. These are designed to provide a sequenced pathway across three implementation levels that builds upon each level:

- Implementation level 1 (IL1) focusses on developing trust and detection foundations to disrupt the most observed malicious actor entry and escalation paths.
- Implementation level 2 (IL2) hardens the organisation's critical processes and services and privileged environments.
- Implementation level 3 (IL3) securing the endpoints and internal enterprise applications.
- Embedded across all levels are the protect surfaces for governance and culture and visibility, which remain enduring priorities.

The measurement model provides a structure for assessment, which includes a four-tier assessment scale that focusses on effectiveness, rather than deployment; not assessed, not effective, partially effective and effective. To support the consistent application of this model, ASD will publish an updated assessment guide as a companion to this framework. This includes providing examples of attestations for each effectiveness scale. An assessor must consider the intersection of the three layers (the cyber security principle, the implementation approach and the protect surface) to determine what an effective implementation should look like. This three-way intersection is what makes the evolved framework different from a checklist.

Part 1: Purpose

The Australian Signals Directorate (ASD) provides cyber security advice and services to government, critical infrastructure, industry, small businesses and individuals. As part of this function, ASD developed the Essential Eight technical guidance which was designed for the protection of enterprise technology networks. Since the initial release, emerging technologies and global cyber threats have evolved. ASD is updating its cyber security guidance to provide contemporary guidance to support organisations to protect themselves against various cyber threats.

With this outlook in mind, ASD is shifting to a broader, principles-based 'Essentials' series. As part of this series, ASD is evolving the existing Essential Eight to the Essentials for enterprise IT – with subsequent publications for other technology environments to follow, including focuses on cloud and operational technology (OT). Information on subsequent publications will be forthcoming.

This consultation paper presents the proposed Essentials for enterprise IT which is designed for all organisations, particularly for Commonwealth entities and critical infrastructure (CI) owners and operators. This represents an evolution of the Essential Eight, redesigned to provide a principles-based approach to cyber security, while offering flexibility in implementation and maintaining a clear and measurable path towards an improved security posture.

The framework evolution considers these key design requirements:

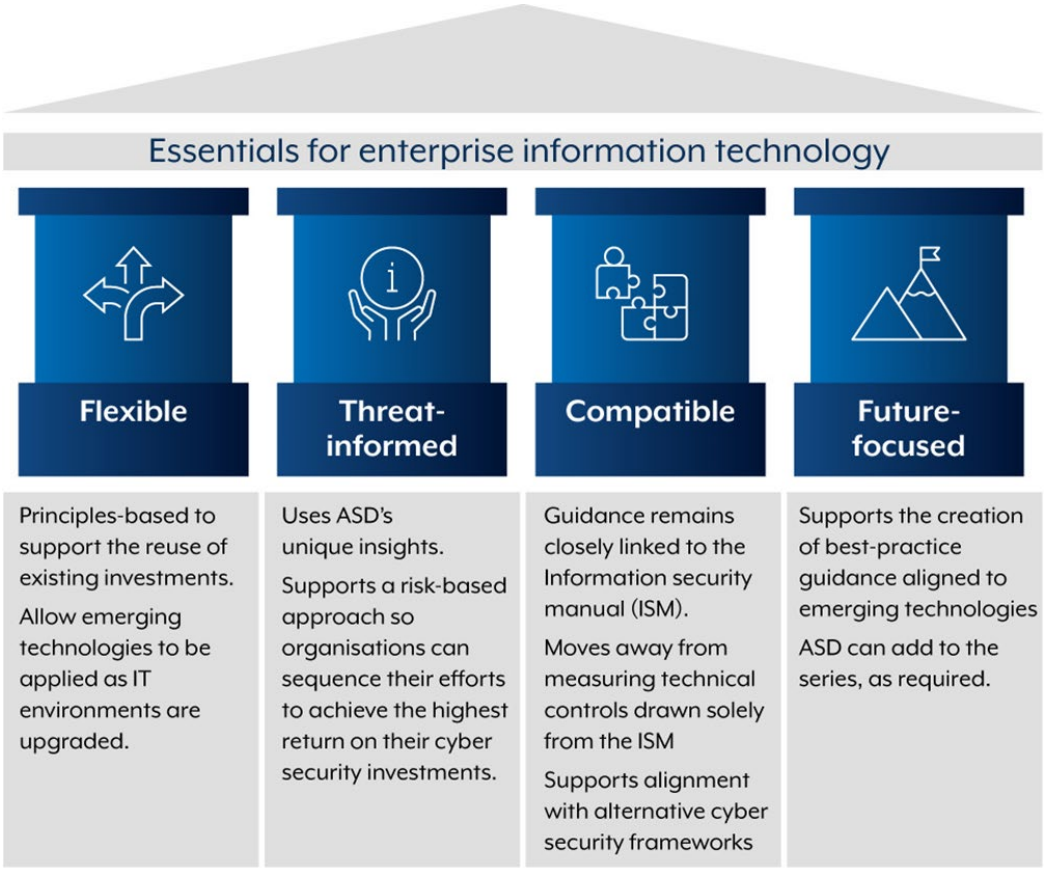


Figure 1: Essentials design principles

The consultation paper is provided to you as an important stakeholder, for you to review the evolved framework's structure, sequencing model and usability before finalisation. Your feedback will help to shape the proposed Essentials for enterprise IT, to ensure that it accurately reflects the current and future needs for Australia's cyber security posture. Please provide your feedback by making a written submission via the ASD Partner Portal. Submissions close on 12 July 2026.

How to read this document

This document is written for multiple audiences involved in cyber security governance, implementation and assessment. Different sections may be more relevant depending on your role. Additionally, comprehensive details is provided to assist readers to understand the need for change, the key concepts and implementation process. Some information presented will fall under supplementary guidance to support implementation.

Audience	Relevant sections (linked)
Chief information security officer Executive or security leader	The evolving cyber threat environment What has changed Framework structure Implementation levels Reporting posture
Security architect Cyber security program lead	Protect surfaces Implementing the framework Getting started Measuring implementation effectiveness
Security assessor or auditor	Measuring implementation effectiveness Assessment process guidance Annex A
Security engineer IT professional Implementation team	Protect surfaces Implementation approaches Annex A
Board member or senior executive	What has changed A different posture for different cyber threat environment Reporting posture

The evolving cyber threat environment

Australia's dependency on digital and internet-connected technology means we remain an attractive target for state-sponsored cyber threat actors, who remain an active and dynamic threat to our nation. The speed of emerging technologies, such as frontier artificial intelligence technology, presents opportunities and emerging threats to the cyber landscape. ASD's Annual Cyber Threat Report 2024-25 highlights that state-sponsored cyber threat actors continue to pose a serious and growing threat to Australia. The report also communicated that state-sponsored cyber threat actors routinely target Australia's networks, possibly to conduct espionage or to preposition for disruptive and destructive effects in the event of a crisis or conflict. The report also highlights that cyber threats affect organisations across all sectors, regardless of size or profile. For Commonwealth entities in particular, ASD's Commonwealth Cyber Posture Report 2025 notes that Commonwealth entities are a common target for malicious cyber activity.

Given the cyber threat environment and technology landscape have evolved significantly since the Essential Eight was released, ASD is updating its guidance to reflect contemporary tradecraft and modern enterprise IT environments. This is to ensure the advice remains effective, achievable and defensible. ASD recommends that organisations continually review their risk and security settings, to protect themselves from the dynamic cyber security landscape. ASD will do its part to provide advice and guidance to mitigate these threats.

The proposed changes have been informed by ASD's unique intelligence insights combined with cyber security incident data, partner feedback and public cyber threat reporting. ASD has taken a principled approach to offer prioritised, threat-informed mitigations for modern IT environments, supported by clear progression pathways and improved usability. This provides the strong basis to evolve ASD's technical guidance, to support Australia to adapt to evolving cyber threats.

What is enterprise IT

Enterprise IT refers to the information technology environments operated by organisations whose scale, complexity, regulatory exposure or reliance on technology requires a structured and deliberate approach to the management of cyber security. Enterprise IT environments typically include multiple interconnected systems supporting different business functions, identity and access arrangements that span a workforce and an extended user population, dedicated cyber security and IT functions with defined roles, and formal governance over technology investment, risk and compliance.

What has changed

ASD recognises that organisations will need time and support to implement the proposed changes. This includes reviewing associated programs for training, assessors and reporting. We are working closely with the various regulatory bodies, who will consider whether future regulatory changes, including reporting, for which transitional arrangements will be required.

The Essentials for enterprise IT has been redesigned around five structural shifts:

Prescriptive controls to flexible principles: The Essential Eight defined specific controls at each maturity level. The Essentials for enterprise IT is built on cyber security principles drawn from ASD's ISM. Each cyber security principle is supported by 'implementation approaches' that describe what needs to be achieved without prescribing a specific technology or product. This makes the evolved framework applicable across diverse IT environments, from on-premises infrastructure to hybrid-cloud architectures. This flexibility also supports enhanced productivity, as some Essential Eight controls were not suitable for all operating environments.
Fixed maturity to targeted implementation: The previous maturity model (ML0–ML3) assessed an organisation's entire IT environment against a single level. The evolved framework introduces 'protect surfaces', which are logical groupings of enterprise functions and assets. This allows organisations to focus their hardening efforts where they will have the greatest impact against malicious actors targeting specific functions.
Uniform progression to sequenced implementation: Rather than advancing all controls uniformly, the evolved framework introduces three implementation levels (IL1-IL3) that recommend a sequence of which protect surfaces should be addressed first. This reflects the reality that not all parts of an IT environment carry equal risk and that resources should be directed to the areas of highest risk. This allows organisations to focus their efforts and resources on protecting business critical functions, rather than an all-or-nothing approach. The implementation levels do not represent assurance over the effectiveness of controls, compliance against a standard, or the cyber security maturity against a particular type of adversary.
Maintaining a broad threat approach, with increased focus on the level of sophistication of cyber threats: The Essential Eight maturity model aimed to mitigate increasing levels of tradecraft associated with malicious actors, where organisations where required to consider the likelihood of being targeted by these actors. In contrast, the *Essentials for enterprise IT* is based upon the state-sponsored threat modelling across all implementation levels. ASD recommends that all organisations adopt this framework to protect themselves against dynamic cyber threats. Organisations that handle sensitive data, support critical functions or form part of the cyber supply chains servicing government or CI may also be targeted. Malicious actors frequently target the most vulnerable organisations in a cyber supply chain to reach the ultimate holder of the data they seek.
Tailored for Commonwealth entities and CI operators, but remains relevant for all organisations: The Essential Eight was designed to provide guidance for organisations ranging from small-to-medium businesses through to large enterprises (including Commonwealth entities). While the Essentials for enterprise IT **remains relevant to organisations of varying sizes and operating environments**, it is tailored for Commonwealth entities and CI operators as their risk profile means they are likely to be targeted by state-sponsored actors.

A different posture for a different cyber threat environment

The previous Essential Eight required Commonwealth entities and most CI operators to meet maturity level two (ML2). This reflected a regulatory assessment that ML2 represented a risk accepted position of the cyber threat environment at the time. ML3 was reserved for organisations considered to be at risk of targeting by state-sponsored actors.

The cyber threat environment has evolved since the release of the Essential Eight. Sophisticated tradecraft is now routinely used against a much broader set of Australian organisations. State-

sponsored actors operate at scale and target vulnerable organisations as steps towards higher-value compromises through their cyber supply chains and trusted relationships. The set of organisations realistically exposed to capable state-sponsored activity has grown well beyond those traditionally associated with the highest level of risk.

The Essentials for enterprise IT reflect this shift. The evolved framework presents an iterative, sequenced path that organisations are expected to progress through and continue to mature. Implementation should not be considered a one-time activity that concludes at the end of the framework. Organisations are expected to maintain, reassess and refine the controls applied at each level on an ongoing basis as the threat environment, the organisations operating context and the framework itself continue to evolve.

Implementation level 3 represents the comprehensive coverage that all organisations should achieve, with the implementation levels providing a sequenced approach to getting there. The evolved framework does not expect that organisations remain at an interim implementation level indefinitely. The intent is that organisations work through the implementation levels at a pace appropriate to their current capability, resourcing and risk environment, with implementation level 3 representing the destination rather than an aspirational ceiling.

ASD recognises that organisational size alone does not determine cyber risks, and that the threat facing Australian organisations affect how businesses of all sizes should approach their cyber security posture. Smaller organisations should consider adopting the guidance in this publication particularly where they operate enterprise technology, form part of government of CI supply chains, or hold sensitive data, as they face threat comparable to those targeting larger entities.

ASD publishes simplified and practical guidance for smaller businesses, home offices and individuals through the small business cyber security resources. The cyber security principles described in this publication remain relevant to smaller business environments, and the security outcomes they describe are equally important. What differs is how those outcomes are achieved. Some responsibilities are met by service providers and vendors on the organisation's behalf, and the scale and structure of implementation reflects the size of the environment, rather than any change in the underlying intent.

Effort distribution across the implementation levels

The implementation levels are designed to mark progression through the evolved framework for organisations at all stages of their cyber security journey. By sequencing the work across three implementation levels, the evolved framework provides a clear and structured path that organisations can follow with confidence, knowing that investment at each implementation level builds towards a coherent cyber security outcome. Implementation level 1 lets organisations show measurable progress within a realistic timeframe, building confidence and momentum at the start of an implementation program. This early progress directs investment to the protect surfaces that disrupt the most observed malicious actor entry and escalation paths, achieving meaningful risk reduction with a contained scope of work. It also allows organisations to develop the operational practices, governance arrangements and assessment capabilities they will need at implementation level 2 and implementation level 3 while operating against a smaller set of protect surfaces.

Implementation level 2 and implementation level 3 each represent a meaningful step forward in capability and coverage. Implementation level 2 extends hardening to the next highest-consequence assets and the elevated access pathways used to administer them, building on the trust and detection foundations established at implementation level 1. Implementation Level 3 extends coverage to the rest of the enterprise IT environment, reinforced by mature governance and visibility capability to support oversight and detection at scale.

This staged approach is one of the evolved framework's strengths. It recognises that building durable cyber security capability is a process that benefits from sequencing, not a single event. Organisations that plan their implementation programs around the effort needed at each implementation level are well placed to maintain momentum through to implementation level 3 and reach an end state they can sustain over time. A sustainable end state is one where all eight protect surfaces have been addressed, the principles applicable to each are operating effectively, and the organisation has the governance, resourcing and operational capability to maintain that posture as the environment and threat landscape continue to evolve.

Questions for you to consider:

- Are these changes reasonable and logical?
- How will these changes impact your entity or stakeholders?
- What are the most useful ways that ASD could support your organisation to communicate the reasons for change to your executive/board?

Part 2: Framework structure

The evolved framework is composed of three layers:

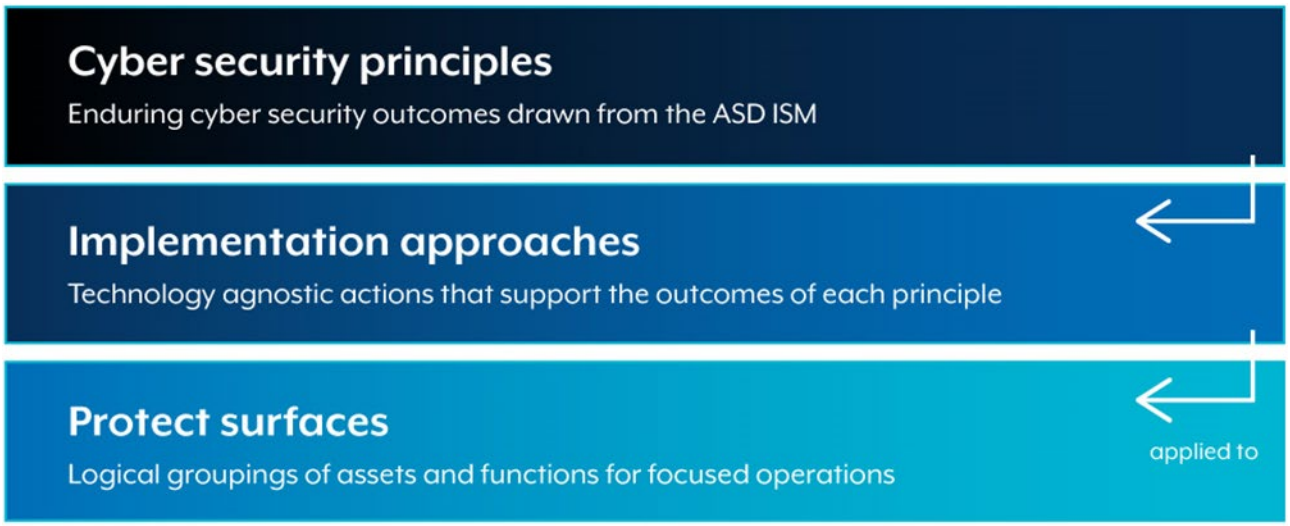


Figure 2: Essentials framework layers

These layers work together to give organisations a clear picture of what to secure, what good looks like and how to get there. Examples are provided throughout the consultation paper.

Cyber security principles

Cyber security principles are drawn from the ISM and represent enduring security objectives for organisations. They describe security outcomes organisations should achieve, rather than the specific mechanisms used to achieve them. Within the Essentials for enterprise IT, organisations are encouraged to prioritise the implementation of a subset of each cyber security principle's scope, as articulated by associated implementation approaches.

The following table includes the list of ISM cyber security principles covered in the Essentials for Enterprise IT.

Table 1: Essentials for Enterprise IT – ISM Principles

Cyber security principle	Principle description
Executive cyber security accountability	The board of directors or executive committee is accountable for cyber security.
Executive artificial intelligence accountability	The board of directors or executive committee is accountable for ensuring that artificial intelligence is secure, controllable, human-supervised and used in an ethical and accountable manner.
Cyber security resourcing	Suitable and sufficient personnel and resources are identified, acquired and maintained in support of cyber security activities.
Security risk management assurance	Security risk management activities for an organisation and their systems (infrastructure, operating systems, applications and data) are embedded into organisational risk management frameworks and subject to ongoing monitoring and assurance activities by the board of directors or executive committee.
Supplier cyber security assurance	Systems (infrastructure, operating systems, applications and data) are delivered and supported by trustworthy suppliers whose cyber security practices are independently verified, or otherwise risk-assessed, on a regular basis.

Cyber security principle	Principle description
Legacy system management	Systems (infrastructure, operating systems, applications and data) that cannot meet the cyber security requirements are managed using compensating controls, along with enhanced monitoring and assurance activities, to maintain an acceptable level of residual risk until they can be decommissioned or replaced.
Asset identification	Systems (infrastructure, operating systems, applications, identities, credentials and data) are continually and centrally identified and documented.
Security risk identification	Security risks for an organisation and their systems (infrastructure, operating systems, applications and data) are identified by using current strategic and sector-specific cyber threat intelligence and threat modelling and are documented with associated risk management decisions.
Secure system lifecycle	Systems (infrastructure, operating systems and applications) are planned, designed, developed, tested, deployed, maintained and decommissioned according to their business criticality ratings and security and resilience requirements using Secure by Design and Secure by Default principles and practices.
Cyber supply chain security	Cyber supply chains supporting systems (infrastructure, operating systems, applications and data) are secure, resilient and support effective and coordinated cyber security incident response.
Identity, credential and access management	Robust and secure identity, credential and access management is used to establish, maintain and control access to systems (infrastructure, operating systems, applications and data) and to support effective detection of identity and credential misuse.
Least privilege access	Personnel and services are granted the minimum access to systems (infrastructure, operating systems, applications and data) required to undertake their duties.
Secure administration	Systems (infrastructure, operating systems, applications and data) are administered in a secure, accountable and auditable manner.
Secure configuration management	Systems (infrastructure, operating systems and applications) are securely configured to approved and maintained baselines, including by reducing attack surfaces and attack paths, with configurations continually monitored and consistently enforced.

Cyber security principle	Principle description
Vulnerability management	Vulnerabilities in systems (infrastructure, operating systems, applications and data) are identified, documented, validated and prioritised for remediation or mitigation in a timely manner, with all remediation and mitigation actions verified for effectiveness.
Trustworthy software	Systems (operating systems and applications) only permit the execution of software that is supported, verified and authorised.
Regular and proven backups	Systems (infrastructure, operating systems, applications and data) are backed up in a secure and proven manner on a regular basis, including through the validation of restoration capabilities.
Network segmentation and segregation	Systems (infrastructure, operating systems and applications) are segmented into network zones based on business criticality and trust levels, with only authorised, controlled and monitored inbound and outbound communication paths between network zones permitted.
Content filtering	Data communicated between different security domains for systems (infrastructure and applications) is controlled and subject to inspection and verification.
Cyber security awareness training	Personnel are provided with ongoing cyber security awareness training, including operational security considerations, tailored to their duties, access levels and current cyber threats.
Centralised event logging	Security-relevant event logs and configuration changes for systems (infrastructure, operating systems, applications and data) are centrally collected, protected against unauthorised modification or deletion, and retained to support effective cyber security event detection and investigation.
Baselined high-risk access activities	Baseline patterns of identity and credential access activities, privileged access activities, and remote access activities are established and maintained for systems (infrastructure, operating systems, applications and data) to enable the detection of anomalous or unexpected behaviour.
Cyber security event detection	Anomalous or unexpected events and behaviours for systems (infrastructure, operating systems, applications and data) are analysed in a timely manner to detect cyber security events.

Cyber security principle	Principle description
Cyber security incident identification	Cyber security events are analysed in a timely manner to identify cyber security incidents.
Detection capability efficacy	Cyber security event detection capabilities are regularly evaluated for effectiveness and continuously refined, including by using current strategic and sector-specific cyber threat intelligence, to improve the detection of cyber security events.
Cyber security incident coordination	Cyber security incident roles and responsibilities are defined, documented and exercised to support the internal and external coordination and management of cyber security incidents, including responsibilities for declaring cyber security incidents and undertaking preapproved response and recovery activities.
Cyber security incident response	Cyber security incidents are contained, eradicated and recovered in a timely manner.
Cyber security incident insights	Lessons learned from cyber security incidents are captured, and areas for improvement are identified, prioritised and actioned in a timely manner.

Implementation approaches

Each cyber security principle from ASD's ISM is tailored through implementation approaches that describe the prioritised administrative or technical outcomes required for its implementation. Implementation approaches are designed to be scalable and technology-agnostic, meaning they describe what must be done without mandating a specific vendor, product or architecture. Implementation is designed to prioritise and deliver the high-impact investment, to achieve the desired security outcomes. The implementation approaches can be found in Annex A, alongside the mapping to the relevant protect surfaces.

Protect surfaces

A protect surface is a logical grouping of enterprise functions and associated critical assets that share a common risk profile. Rather than attempting to secure the entire IT environment uniformly, the evolved framework directs organisations to identify and harden discrete protect surfaces in a deliberate sequence. The evolved framework defines Eight protect surfaces, each assigned to an implementation level that reflects its priority in the sequencing model.

Implementation Level	Protect surface	Description
All	Governance and culture	The policies, risk management processes, roles, responsibilities and security culture that underpin the cyber security program.
All	Visibility and response	The procedures and capabilities to detect, analyse and respond to cyber security events and incidents.
1	Boundary services	Systems positioned at network boundaries that broker connectivity between internal and external environments, including internet-facing services and network edge devices.
1	Identity services	Identity, credential and access management services including directories, authentication and authorisation enforcement solutions and credential stores.
2	Critical processes and services	Core processes and services that are essential to delivering primary organisational outcomes or contributing to personnel health and safety.
2	Privileged environments	Systems and services that support privileged activities, management services and high-impact administrative functions.

Implementation Level	Protect surface	Description
3	Business systems	Supporting enterprise functions that collectively enable organisational operations.
3	User devices	Endpoints issued to the workforce, including desktops and laptops.

Implementation levels

Implementation levels define the recommended sequencing for addressing protect surfaces. While the implementation levels are not traditional maturity ratings, they do reflect the maturing resilience of an organisation against malicious actor targeting. As an organisation progresses through the levels, it extends controls to a wider range of protect surfaces and reduces the number of viable pathways a malicious actor can use to achieve their objectives.

The implementation model is designed to support entities to move through implementation levels 1 through to 3. Previously, entities would select a maturity level (ML0-ML3) to mitigate tradecraft associated with different levels of adversary tradecraft. The revised approach for measuring implementation levels, provides clearer indications of where organisations are against their security goals, rather than against static tradecraft.

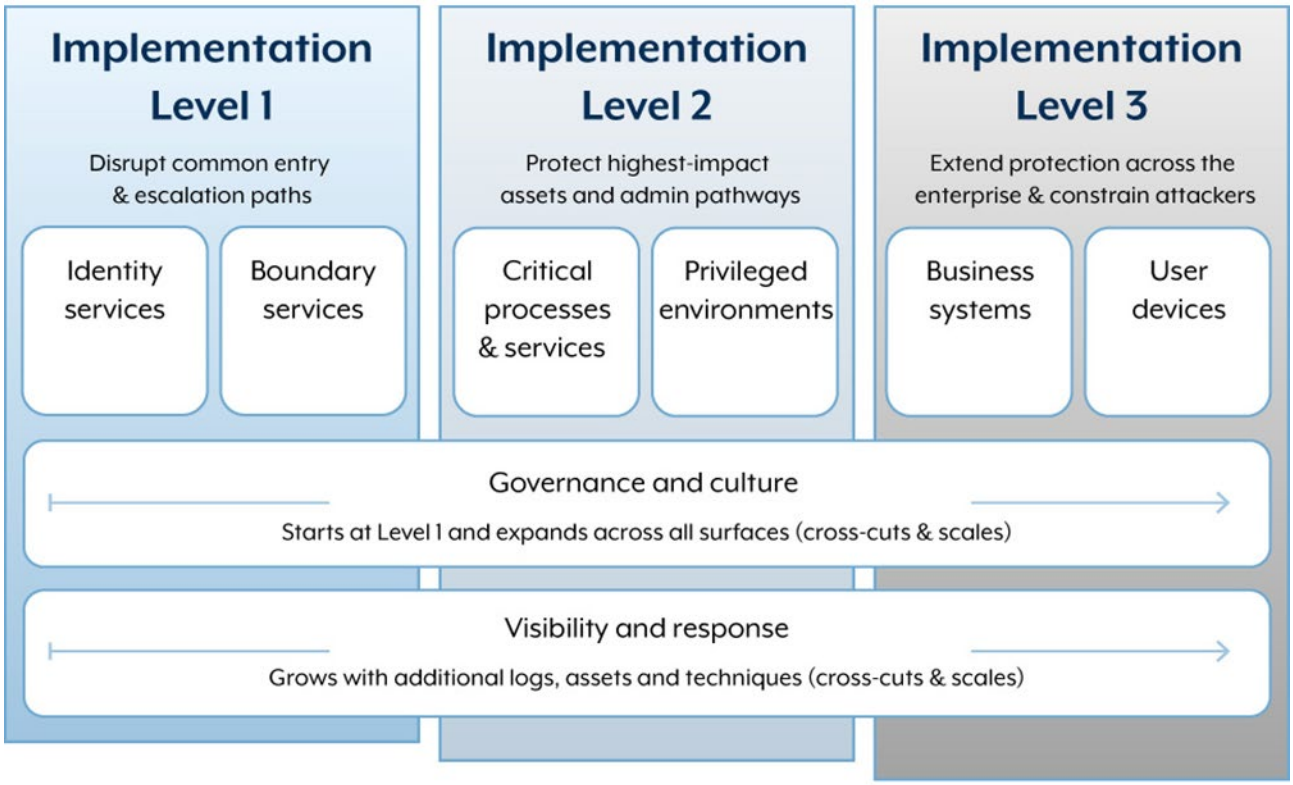


Figure 3: Essentials implementation levels

Implementation level 1

Implementation level 1 (IL1) addresses 'boundary services' and 'identity services'. These two surfaces represent the entry point and the trust anchor of the environment. 'Boundary services' controls connectivity and data flows between environments of different trust levels, including the internet-facing surface that malicious actors most commonly target for initial access. 'Identity services' secure the directory, authentication and authorisation mechanisms solutions that control access across the enterprise IT environment which are common targets of malicious actors. Hardening these two surfaces first directly disrupts the most frequently observed malicious actor entry and escalation paths.

IL1 marks the point at which organisations begin establishing 'governance and culture' and 'visibility and response'. Unlike the other protect surfaces, these two protect surfaces are not bound to a single implementation level. They are foundational and continuous protect surfaces that apply across all three levels. 'governance and culture' ensure that decisions, policies, risk ownership, and training are aligned to deliver the organisation's desired security outcomes. 'Visibility and response' provide the detection, analysis and response capability that identifies and contains malicious actor activity wherever it occurs in the environment. Both surfaces should be considered at every stage of implementation. Organisations should begin addressing them at implementation level 1 and continue to mature them as work progresses through levels 2 and 3, recognising that the requirements placed on these surfaces evolve as additional protect surfaces come into scope. Without these foundations, investment in later protect surfaces is unlikely to deliver lasting or effective improvements.

Implementation level 2

Implementation level 2 (IL2) addresses 'critical processes and services' and 'privileged environments'. These two surfaces represent the next-highest-consequence assets in the environment and the elevated access pathways used to administer them. 'Critical processes and services' protect the systems and information whose compromise would cause the greatest organisational impact. 'Privileged environments' secure the dedicated infrastructure, platforms and tooling used to perform high-impact administrative functions. Addressing these surfaces at IL2 builds on the controls established at IL1 and directs deliberate hardening to the assets and pathways most aggressively targeted by capable malicious actors.

'Governance and culture' and 'visibility and response' continue to evolve at this level, with governance arrangements extending to cover the additional surfaces in scope and detection capability extending to the additional log sources, asset types and malicious actor techniques relevant to those surfaces.

Implementation level 3

Implementation level 3 (IL3) addresses 'business systems' and 'user devices'. These two surfaces represent the broader enterprise environment that supports day-to-day operations and the endpoints through which the workforce interacts with that environment. 'Business systems' cover

internal enterprise applications, collaboration platforms and supporting infrastructure. 'User devices' address the endpoints issued to or used by personnel.

These surfaces are sequenced at IL3 because the utility of compromising them is significantly reduced when the principles and implementation approaches applied at IL1 and IL2 are in place. An adversary who reaches a user device or a business system in an environment with strong identity, hardened boundaries, protected critical processes and disciplined privileged environment has limited opportunity to escalate privileges, move laterally to higher value targets or extract significant value from the foothold they have established. Hardening these surfaces is still necessary to reduce the residual risk and to complete the frameworks coverage of the enterprise, but the earlier levels deliver the controls that constrain what an adversary can do once they arrive.

'Governance and culture' and 'visibility and response' continue to mature at this level, extending governance coverage to the full enterprise environment and ensuring that detection and response capability addresses the full scope of systems and user activity.

Questions for you to consider:

- Does the framework present a logical flow to the structure?
- Does this shift from prescriptive controls to a principle-based approach create any challenges or benefits for your organisation?
- Is it clear that the Implementation Levels differ in intent to the Essential Eight maturity levels? If not, would a different term provide clarity?

Understanding the protect surfaces

A protect surface is a logical grouping of enterprise functions and associated assets that share a common risk profile. The purpose of scoping is to identify which critical assets within organisations fall within each protect surface such that the applicable cyber security principles and associated implementation approaches can be implemented in a targeted manner.

When scoping, organisations need to consider assets across these categories:

- *Infrastructure:* The underlying hardware, software and networking components that are required to run IT within an organization, including physical and virtual servers, and network systems such as switches and routers
- *Operating systems:* System software that manages hardware and software resources, including Microsoft Windows, Apple macOS and Linux distributions
- *Applications:* Software programs designed for end users, including office productivity suites, web browsers, or email clients
- *Data:* Information and data assets that are important for business operations and systems, including user objects, configuration information, and other data sources

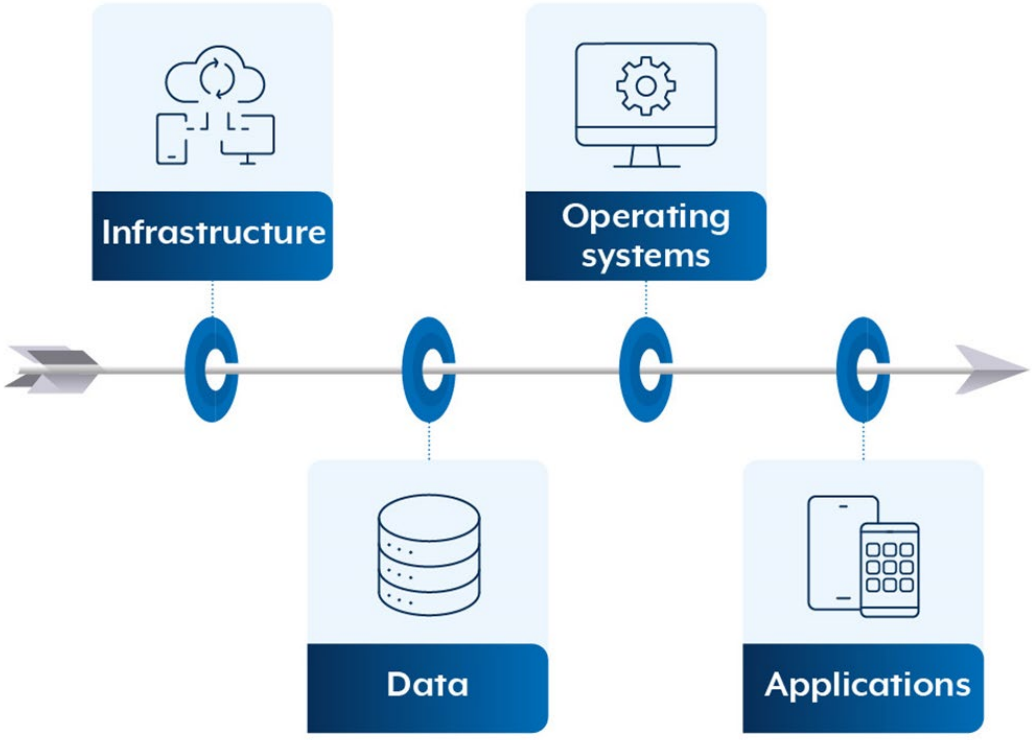


Figure 4: Protect surface asset categories

Not every protect surface will have significant assets in every category.

The following pages describe the eight Protect Surfaces and provide examples of the assets, programs and artefacts in scope for each one.

Governance and culture

What this protect surface involves

Governance and culture represent the administrative foundations that give the cyber security program its authority and longevity, including the policies, risk management processes, roles, responsibilities and security culture that underpin every other control. This surface is primarily people focused and process-driven rather than technology-driven.

Consequence of compromise

Weak governance and cyber security culture can result in inconsistent security decisions, unclear accountability, delayed remediation activity and reduced organisational ability to manage cyber security risk over time. Deficiencies in this protect surface may undermine the effectiveness and sustainability of controls implemented across all other protect surfaces.

Examples of assets in scope

- overarching organisational cyber security strategy documents, including cyber security policies and procedures
- security awareness and training programs
- employee code of conduct policies
- security roadmaps
- risk registers

Visibility and response

What this protect surface involves

Visibility and response represent the organisation's ability to detect, analyse and respond to cyber security events and incidents. This protect surface encompasses the procedures, platforms and capabilities that collect security telemetry then convert it into actionable insight to enable coordinated response and recovery activities when incidents occur.

Consequence of compromise

Without effective visibility, malicious actors can maintain persistence, escalate privileges and cause impact without opposition. This can slow incident identification and response, leading to prolonged impact, higher resources to resolve, and increased organisational harm.

Examples of assets in scope

- asset registers
- incident response plans and procedures

- endpoint detection and response solutions
- extended detection and response solutions
- security incident and Event management solutions

Boundary services

What this protect surface involves

Boundary services represent the systems and services that control connectivity and data flows between environments of different trust levels, including but not limited to internet-facing services, cross-domain gateways, interagency connections, cloud service boundaries and OT boundaries. This protect surface is often the organisation's front door and the most visible target for external malicious actors.

Consequence of compromise

Malicious actors routinely scan for and exploit vulnerabilities in internet-facing services as a common method of enabling initial access. Email gateways, virtual private network (VPN) concentrators, web applications and remote access portals are among the most frequently targeted assets. A compromised boundary service can enable initial network access, data interception, service disruption visible to customers and partners, and staging of attacks against third parties.

Examples of assets in scope

- VPN concentrators
- remote access solutions
- enterprise routers and switches that facilitate traffic from the internet to internal networks
- web servers and web applications
- application programming interface endpoints
- software-as-a-service (SaaS)
- email gateways
- network and web application firewalls

Identity services

What this protect surface involves

Identity services represent the functions responsible for authentication, authorisation, identity lifecycle management and access control policy enforcement. This surface is the trust anchor for

every other technical protect surface. If malicious actors control identity, they typically control access to everything.

Consequence of compromise

Compromise of identity services enables broad, cascading access across the IT environment. Malicious actors who gain control of directory services, federation platforms or credential stores can impersonate any user, escalate to any privilege level and move laterally without triggering controls that rely on identity as a trust signal. Compromise of identity services can occur at multiple stages of an attack chain, from initial credential theft to the establishment of persistent attacker-controlled access that survives password resets, account disablement and even partial environment rebuilds.

Examples of assets in scope

- identity, credential and access management solutions
- directory services
- federation services
- certificate services

Critical processes and services

What this protect surface involves

Critical processes and services represent the core services essential to delivering primary organisational outcomes or contributing to personnel health and safety. This is the protect surface where compromise is most likely to translate directly into mission or business objective failure.

Consequence of compromise

Malicious actors targeting this protect surface typically pursue disruption (rendering services unavailable through ransomware or destructive attacks) or theft (exfiltration of high-value data for intelligence, competitive advantage or extortion). Because these assets represent the highest-value targets, they attract the most capable and deliberate malicious actor activity.

Examples of assets in scope

Examples of assets that may fit within this protect surface include:

- intellectual property data stores
- customer data, personally identifiable and personal sensitive information data stores
- financial or payment services and capabilities
- key business applications and services

- email systems and other key communication platforms

Privileged environments

What this protect surface involves

Privileged environments represent the dedicated infrastructure, platforms, tooling, accounts and controls that enable and govern privileged activities and high-impact administrative functions. This protect surface encompasses the environments from which administrative actions are performed, the platforms that broker and control privileged access, and the policies and mechanisms that constrain how elevated permissions are granted, used and recorded.

Consequence of compromise

Malicious actors who obtain privileged access can modify configurations, disable controls, create persistence mechanisms and access any data within scope of those privileges. Techniques such as credential theft, token manipulation and exploitation of misconfigured delegation are standard components of post-compromise activity. Once privileged environments are compromised, malicious actors operate with the same authority as the organisation's own privileged users, making detection significantly more difficult.

Examples of assets in scope

Examples of assets that may fit within this protect surface include:

- privileged user accounts
- secure admin workstations and privileged access workstations
- privileged identity management solutions
- privileged access management solutions
- jump servers
- cloud management platforms

Business systems

What this protect surface involves

Business systems represent the supporting enterprise functions that enable day-to-day operations. These are not individually critical, but their collective compromise degrades the organisation's ability to function.

Consequence of compromise

Malicious actors frequently use business systems as staging points to move laterally towards higher-value targets. The disruption of multiple business systems in parallel can render the organisation unable to perform coordination, communication and logistical activities.

Examples of assets in scope

Examples of assets that may fit within this protect surface include:

- corporate management systems
- cloud applications
- IT service management (ITSM)/IT infrastructure library (ITIL) systems
- office productivity applications
- collaboration and messaging tools
- data repositories

User devices

What this protect surface involves

User devices represent the endpoints issued to the workforce. This protect surface is the most distributed and the most exposed to user-initiated risk.

Consequence of compromise

User devices are the primary vector through which phishing, social engineering and drive-by compromises reach organisations. A compromised user device gives malicious actors access to a foothold for lateral movement. User devices are also the most common point from which data leaves the organisation; whether through exfiltration, accidental exposure or unsanctioned applications. The volume and diversity of user devices make this protect surface the most challenging to secure consistently.

Examples of assets in scope

Examples of assets that may fit within this protect surface include:

- desktop computers and laptops
- virtual desktops for standard user access
- enterprise managed mobile devices, such as mobile phones and tablets

Functions that span multiple protect surfaces

Not every function or asset will fall cleanly within a single protect surface. Enterprise functions may span more than one surface, either because the function brings together different categories

of assets or because the same asset services multiple roles within the environment. The framework accommodates this through a simple rule: where an asset or function maps to more than on protect surface, the implementation approaches applicable at the earliest implementation level should be applied first, with the approaches from later levels added as the organisation progresses through the framework.

Example: A critical internet-exposed application

An organisation operates a customer-facing application that processes sensitive personal information. The application is critical to the organisation's primary outcomes (critical processes and services) and is exposed to the internet through a web boundary (boundary services). Under the rule described above, hardening begins at IL1 because the application fits into the definition of boundary services. The organisation applies the boundary services implementation approaches to the applications internet-facing components, including ingress controls, web application firewall configuration, inbound content inspection and boundary device hardening. When the organisation moves to IL2, the critical processes and services implementation approaches are applied, where the scope of assets or approaches is different to those in IL1. The application is now hardened against both boundary-facing and mission-critical threats

Questions for you to consider:

- Are there any areas where Protect Surfaces boundaries are unclear?
- Do the protect surfaces effectively cover your organisation's enterprise environment?
- What additional information or guidance would you need to assist with scoping activities?

Part 3: Implementing the framework

The framework applies principles and implementation approaches to protect surfaces through a defined mapping (Appendix A refers). Not every principle and implementation approach applies to every protect surface. Principles or implementation approaches are not applied where they do not meaningfully reduce risk for the assets within a surface, or where the relevant risk is more effectively addressed through other principles.

Some principles apply more broadly across the enterprise IT environment, while others are more focused on targeted areas. Principles such as vulnerability management, secure configuration management and centralised event logging apply to a broader range of the protect surfaces that are more vulnerable to software exploitation. Principles from the 'govern' function apply primarily to the 'governance and culture' protect surface, and principles from the 'detect' and 'response' functions apply primarily to 'visibility and response' protect surface. Some principles, such as 'network segmentation and segregation', apply only where the underlying control is meaningful for the assets in scope for the protect surface.

Options for implementation focus

The evolved framework supports two flexible options for implementation.

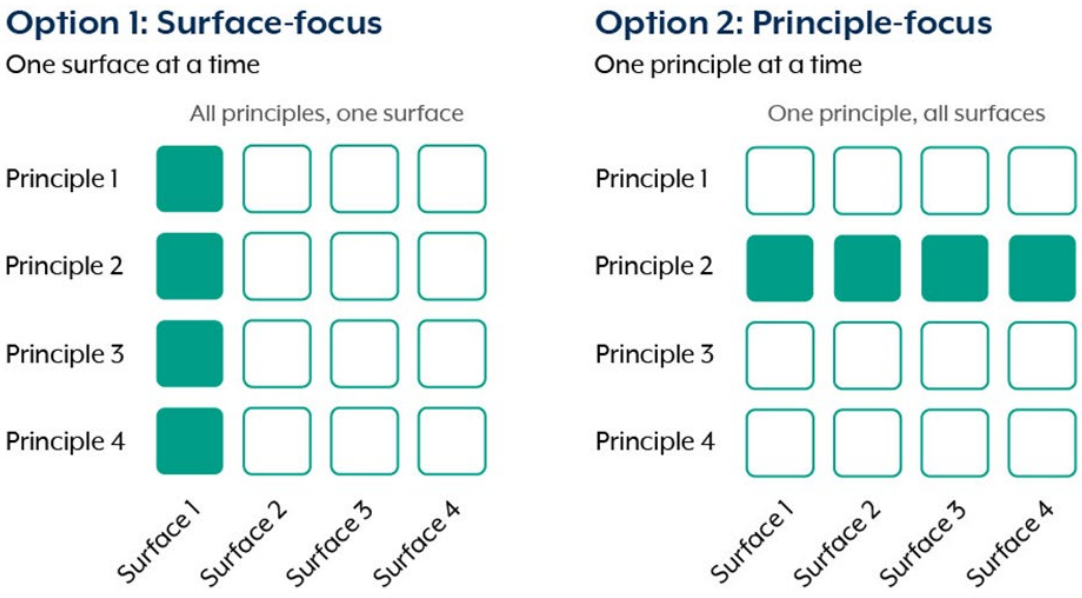


Figure 5: Framework implementation option

Option 1 - Surface-focus

In a surface-focus approach, organisations select a single protect surface and implement all applicable implementation approaches for each cyber security principle before moving to the next protect surface. This approach is well suited to organisations that want to demonstrate measurable hardening of a specific function within a defined timeframe.

The primary benefit of this approach is that it directs efforts to the most critical assets in priority order. Each protect surface within the implementation group are addressed before the organisation moves to the next group, meaning the assets representing the highest risk to the organisation receive comprehensive controls earlier in the implementation program. This approach produces clear, demonstrable outcomes for each protect surface, which is useful for reporting to executives, regulators or boards who want to see tangible progress against a defined scope.

The trade-off to consider with this approach is the impact on licensing and procurement. This is because in a surface-focus approach controls are applied to a single protect surface at a time, whereas the organisations may purchase tooling, licences and services in batches to be applied to multiple protect surfaces as a single enterprise-wide commitment. For some organisations this can represent a cost-management advantage as investment is staged in line with delivery, expenditure can be aligned to budget cycles, and the organisation retains flexibility to adjust scope or vendor selection between surfaces based on lessons learned. Other organisations may find that larger upfront commitments deliver better commercial terms. The right choice depends on the organisation's procurement model, cash flow preferences and confidence in long-term suitability of selected solutions.

Example: Surface-focus approach

An organisation chooses the 'identity services' protect surface. The organisation identifies the relevant cyber security principles, reviews the associated implementation approaches, and applies the applicable implementation approaches to all identity-related assets (e.g. directories, authentication services and credential stores). Once satisfied that all implementation approaches have been met for the protect surface, the organisation moves to the next protect surface.

Option 2 - Principle-focus

In a principle-focus approach, organisations select a cyber security principle and applies its implementation approaches across multiple protect surfaces, using the implementation level sequencing to determine which assets are addressed first. This approach is well suited to organisations deploying a specific capability or solution across the enterprise and wanting to prioritise where it is applied.

The primary benefit of this approach is broader and more consistent coverage of a given control across the environment. Designing a capability once and applying it across multiple surfaces typically produces more uniform implementation, reduces duplication of effort and allows the organisation to take advantage of enterprise-wide licensing and procurement. Tooling and services are often more cost-effective when purchased at scale, and operational support is simpler when a single capability is deployed consistently rather than as multiple surface-specific instances.

The primary drawback is that effort may be expended on assets that do not represent the best risk reduction. By rolling out a control across all applicable surfaces, the organisation may invest in hardening assets in lower-priority protect surfaces before higher-priority protect surfaces have

been fully addressed. Without disciplined application of the implementation level sequencing, a principle-focus approach can result in uniform but suboptimal risk reduction.

Example: Principle-focus approach

An organisation is rolling out a new multi-factor authentication (MFA) capability. It identifies the relevant cyber security principle (identity, credential and access management), its implementation approaches, then applies the applicable implementation approaches to all IL1 protect surfaces, followed by IL2 protect surfaces and finally IL3 protect surfaces.

Which option to choose?

Organisations that are early in their cyber security journey or that do not yet have the internal capability to confidently consider the trade-offs between approaches, should adopt a surface-focus approach as the default. A surface-focus approach provides a clear, sequenced path through the framework that directs efforts to the most critical assets first. It produces measurable progress against well-defined scopes and reduces the risk of investment being spread too thinly across the environment. Organisations using a surface-focus implementation as the default can transition to a blended or principle-focus model as their capability matures and as cross-cutting opportunities become evident.

Organisations with established cyber security programs and the maturity to assess their own operating context may select the approach that best suits their needs or combine elements of both.

Getting started with implementation



Figure 6: Implementation steps

Step 1: Understand your enterprise IT environment and operating context

Before applying the framework, the organisation must consider their operating context. This may include the organisation's regulatory obligations, international partnerships, stakeholder expectations, and the organisation's risk appetite, resources and funds available for cyber uplift.

Next, the organisation must identify the systems and services within the protect surfaces it intends to address. This is not a comprehensive enterprise-wide asset discovery exercise. The objective is to scope the protect surfaces being targeted in the current implementation cycle, drawing on existing asset inventories, system registers, business impact assessments and architecture documentation where these exist.

Organisations following the recommended sequencing need to start by scoping the two IL1 protect surfaces: 'boundary services' and 'identity services', as well as the two foundational

surfaces, 'governance and culture' and 'visibility and response'. Assets in scope for IL2 and IL3 protect surfaces can be identified later as the organisation progresses through the levels. This staged approach reduces the upfront effort required to begin implementation and allows scoping decisions to be informed by lessons learned from earlier protect surfaces. For each protect surface, organisations should identify the specific systems, services and assets that are in scope. The protect surface definitions are intentionally broad to accommodate diverse IT environments, which means organisations must make scoping decisions about what falls within each protect surface. Assets that serve multiple surfaces should be documented with reference to each surface in which they operate, as the applicable implementation approaches may differ.

Where significant gaps in asset visibility are identified during scoping, the organisation should note these for resolution as part of addressing the relevant protect surfaces. The visibility and response protect surface includes implementation approaches specifically designed to support automated asset discovery and inventory.

Example: Scoping

An organisation begins by mapping its enterprise IT environment to the two protect surfaces in IL1. For each protect surface, it identifies critical assets across four categories: infrastructure, operating systems, applications and data, where applicable

For the 'identity services' protect surface, it identifies:

Infrastructure: Domain controllers, federation servers, RADIUS server and hosting network segments.

Operating systems: Windows Server instances underpinning directory and federation services.

Applications: Active Directory, Entra ID tenant, directory synchronisation tool and a certificate authority.

Data: Directory database, user and group objects, credential material, authentication tokens, federation trust metadata, access policies and audit logs.

Where an asset did not fall neatly into a single protect surface, the organisation documented where assets served multiple protect surfaces and subsequently considered this during their implementation.

Step 2: Conduct a gap analysis against current capabilities

For each cyber security principle applicable to the protect surface being targeted, organisations should review the associated implementation approaches and assess them against existing capabilities. The purpose of this step is not to perform a formal measurement, but to identify which implementation approaches are already addressed and which represent gaps requiring action.

For each implementation approach, organisations should determine whether:

- **The implementation approach is already effectively addressed by existing controls.** No further action is required, but organisations should still document how existing controls satisfy the implementation approach for future assessment purposes.
- **The implementation approach is partially addressed by existing controls.** Existing controls provide some coverage, but gaps in scope, consistency or enforcement have been identified. These should be documented with sufficient detail to inform remediation planning.
- **The implementation approach is not addressed by existing controls.** No existing controls satisfies the intent of the implementation approach. This represents new work that must be planned and resourced.

Example: Gap analysis

The organisation reviews the 'identity services' protect surface against the applicable cyber security principles. Under the 'identity, credential and access management' cyber security principle, it assesses five implementation approaches. It finds that phishing-resistant MFA is deployed for all interactive user authentication to directory services management, federation service consoles and identity management platforms (assessed as **effective**), but that default and unused accounts have not been reviewed in over 18 months (assessed as **partially effective** - a process existed but lapsed) and that no automated scanning for exposed credentials in code repositories or collaboration platforms has been implemented (not addressed). The organisation records these findings and moves to the next cyber security principle, until the gap analysis is completed.

Step 3: Determine the implementation approach for remaining work

For each gap identified in Step 2, organisations need to determine whether a protect surface-focus or principle-focus approach, or a combination, is most appropriate. This decision needs to be based on the nature of the identified gap, current priorities and whether a remediation is best delivered as a targeted effort within a single protect surface or as a broader capability rollout.

Where a gap exists in a single protect surface, then a surface-focus approach is typically more efficient. Where the same gap exists across multiple protect surfaces, a principle-focus approach allows organisations to deploy the capability once and apply it in a sequenced manner across all applicable protect surfaces. This is common for cross-cutting capabilities, such as MFA, configuration management or event logging.

Example: Choosing an approach

The organisation identifies two categories of remaining work at IL1. First, several gaps specific to the 'governance and culture' protect surface. This includes that its security risk management methodology has not been formally approved and its cyber security awareness training has not

been updated for modern threats and risks. These are addressed using a surface-focus approach, as the remediation is contained within that protect surface.

Second, it identifies that the 'centralised event logging' cyber security principle is a gap for the 'identity services' protect surface, and the same gap applies to every technical protect surface at IL2 and IL3. Rather than solving this in isolation, the organisation takes a principle-focus approach: it designs the event log forwarding architecture as a repeatable pattern, applies it to the 'identity services' protect surface first, then extends coverage to IL2 and IL3 protect surfaces as it progresses.

Step 4: Implement and measure

Once an implementation option has been chosen, organisations should begin the execution of planned remediation activities within the targeted protect surfaces. As implementation progresses, each implementation approach should be assessed against its applicable protect surface using the outcome-based measurement scale described in Measurement section of this paper.

Each implementation approach should be assessed for coverage, consistency, currency and evidence. An implementation approach that is deployed but not applied to all relevant assets within the protect surface, or that is inconsistently enforced, should be rated **partially effective** rather than **effective**. Where an approach is rated **partially effective**, the specific gaps should be documented and a remediation plan with defined timeframes should be put in place.

A protect surface is considered addressed when all applicable implementation approaches are operating as intended and therefore rated **effective**.

Step 5: Progress through the implementation levels

Once the targeted protect surfaces have been addressed, organisations should repeat the same approach at the next set of protect surfaces

The implementation levels should be followed in sequence as each level builds on the capabilities established in the preceding one. This supports scalability and efficiency during the implementation. Organisations with specific cyber threat intelligence or regulatory drivers may adjust the order of protect surfaces, however, the provided sequence of implementation levels represents ASD's recommended approach.

Reassess and iterate

The evolved framework is designed to support continuous improvement. Organisations should establish a regular cadence for reassessing their protect surfaces against the applicable cyber security principles and associated implementation approaches. Reassessment should consider changes to the IT environment, the threat environment, organisational priorities and the outcomes of any cyber security incidents since the last assessment.

Organisations should consider periodic independent assessment as part of their reassessment cycle. Independent assessors bring fresh judgement and are more likely to identify gaps that internal teams have normalised over time. Independent assessment also strengthens the credibility of posture reporting to regulators and governing bodies, particularly where external assurance is required.

Exceptions, compensating controls and accepted risks should all be reviewed during reassessment. An exception that has expired without review, a compensating control that has degraded or an accepted risks that has not been reconsidered against the current threat environment all represent opportunities for the organisation to either close the gap or formally re-approve the position.

Where reassessment identifies an implementation approach previously rated as **effective** has degraded, organisations should document the gap and initiate remediation activities using the same process described in steps 2 through 4. The measurement model supports this iterative approach by providing a consistent scale that can be applied at each reassessment.

Questions for you to consider:

- Which implementation approach appeals to your organisation, and why?
- What timeframes would your organisation need to meet the implementation level 1 changes?
- What resourcing and funding implications would the proposed changes have for your organisation?
- Does your organisation have the capability to select the appropriate cyber mitigation?
- What additional guidance or supporting materials would you need to move through the implementation levels?

Measuring implementation effectiveness

The evolved framework uses an outcome-based assessment model to measure how effectively organisations have implemented the cyber security principles and their associated implementation approaches. The evolved framework supports both self-assessment and independent assessment and produces results that can track progress through the implementation levels, report posture to regulators and governing bodies, and direct remediation effort.

A shift in how to measure implementation effectiveness

The previous framework assessed against specific, and often prescriptive, controls. Measurement was typically conducted by assessing whether each control had been implemented as specified. This approach had the advantage of clarity, but it encouraged a compliance mindset.

The evolved framework focuses on achieving underlying security outcomes by using a principles-based approach which is measured differently. Because the cyber security principles and associated implementation approaches are expressed in technology-agnostic, outcome-based language, the measurement focuses on whether the intended security outcome is achieved, and not whether a specific product has been deployed or a specific configuration has been applied.

Example: Secure Configuration Management

Under the previous framework, an organisation could demonstrate compliance with user application hardening by showing that browser setting, productivity application configurations and web content controls had been deployed in line with the published guidance. Under the evolved framework, the equivalent assessment asks whether the hardened configurations have been deliberately defined for the organisation's environment and whether they continue to operate as intended over time. The evidence must demonstrate that the baseline has been developed, formally approved and reviewed on a defined cycle, and that drift from the baseline is detected and remediated within defined timeframes. An organisation that adopted a generic baseline at deployment without tailoring it to its environment, and has limited visibility of configuration drift since, is unlikely to be sustaining the intended security outcome regardless of how thorough the original deployment was.

Organisations accustomed to the previous framework can expect three practical changes.

Greater flexibility in how to satisfy an implementation approach

Two organisations may implement the same implementation approach using different products, architectures or processes. Both can be assessed as **effective** if the security outcome is achieved. Assessors must evaluate the overall effect of the implementation approach to achieve the aligned security principle, not necessarily whether it matches a predetermined technical specification.

Established control frameworks remain a valuable reference point in assessing effectiveness. The ISM provides detailed control statements that align to the principles in this framework, and may

be appropriate depending on the organisations regulatory context. Implementers can use control frameworks to identify specific controls that satisfy an implementation approach, and assessors can use them as a reference point when evaluating whether the deployed controls deliver the intended security outcome.

Evidence of effectiveness replaces evidence of deployment

Demonstrating that a control has been deployed is necessary, but organisations must demonstrate that the control is operating as intended across the relevant protect surface; showing coverage, consistency, currency and supporting evidence.

Measurement is contextual to the protect surface

The same implementation approach may be assessed differently across different protect surfaces depending on the assets in scope and the consequence of failure.

Assessment scale

Each implementation approach is assessed against each applicable protect surface using a four-tier outcome-based scale.

Not assessed: The implementation approach has either not yet been assessed, or insufficient visibility of an implementation existed to determine whether there was a meaningful reduction in risk for the protect surface.

Not effective: The implementation approach is absent or implemented in a way that does not meaningfully reduce risk to the protect surface. The intended security outcome is not being achieved.

Partially effective: The implementation approach is in place but has identified gaps in coverage, consistency or enforcement that limits its effectiveness. The security outcome is being achieved in some areas but not reliably across the protect surface. A remediation plan with defined actions and timeframes is expected. Where gaps have not yet been fully articulated or a remediation plan is still being developed, organisations should ensure that any assessment provides sufficient time to identify the gaps and define the path to closure before a final rating is determined.

Effective: The implementation approach is implemented in a way that demonstrably reduces risk to the protect surface. The intended security outcome is being achieved consistently, with supporting evidence available where applicable.

Assessment process guidance

The measurement model described in this section provides the structure for assessment, including the four-tier assessment scale. To support the consistent application of this model, ASD will publish an updated assessment guide as a companion to this framework. The assessment guide will replace the existing *Essential Eight assessment process guide* and is intended for use by both internal assessors performing self-assessment and independent assessors providing external assurance. The Essentials for enterprise IT draw on the assessment approach from the United

Kingdom's National Cyber Security Centre's (NCSC) Cyber Assessment Framework, which utilises *Indicators of Good Practice* to assist entities in measuring principle-based guidance.

The assessment guide will provide Indicators of Good Practice (IGP) for every implementation approach, organised by principle. For each approach, the guide will set out the statements that distinguish a **not effective**, **partially effective** and **effective rating**. The statements are intended to inform expert judgement rather than serve as a checklist. The presence of any single **not effective** statement is normally sufficient to justify a **not effective** rating that approach. A **partially effective** rating must reflect delivery of specific worthwhile security benefits, not merely partial activity against the approach.

The following examples illustrate the format that the assessment guide will use:

Approach: (Executive cyber security accountability) Formally define and establish executive and board-level accountabilities for cyber security

Not effective	Partially effective	Effective
At least one of the following statements is true:	**All of the following statements are true:**	**All of the following statements are true:**
No named executive or board member is accountable for cyber security Cyber security accountability exists only informally, or is assumed rather than documented No evidence of cyber security reporting to the board of executive	A named executive or board member is accountable for cyber security, but reporting obligations are not formally defined Board-level reporting on cyber security occurs but is ad hoc rather than on a defined schedule Accountability is documented but has not been reviewed since organisation restructuring	A named executive or board member has documented accountability for cyber security with defined reporting obligations Regular reporting to the board or executive occurs on a defined cadence with evidence of reporting artefacts Accountability arrangements are reviewed when organisational structures or leadership changes occur

Approach: (Identity, credential and access management) Identify and disable or remove default and unused accounts on a regular basis

Not effective	Partially effective	Effective
At least one of the following statements is true:	**All of the following statements are true:**	**All of the following statements are true:**

Not effective	Partially effective	Effective
No review of default or unused accounts has been conducted within the defined period Default vendor accounts remain enabled on production systems No process exists to identify accounts that have not authenticated within a defined period	A review has been conducted but does not cover all assets within the protect surface Default accounts have been identified but some remain enabled without documented justification Unused accounts are identified but remediation is incomplete or delayed beyond defined timeframes	A review of accounts has been completed with defined review period covering all assets in the protect surface Default accounts are disabled or have documented and approved justification for remaining active Unused accounts are identified by last authentication date and disabled or removed within the defined timeframe.

The assessment guide will also provide guidance on how to interpret the descriptors in the context of each protect surface, what evidence and artefacts assessors should look for, how to handle situations where compensating controls are in place, and how to document assessment findings in a form suitable for reporting to regulators and governing bodies. The guide will be developed in consultation with the assessor community and will draw on the lessons learned from the existing *Essential Eight assessment process guide*.

Assessment across protect surfaces, principles and approaches

The evolved framework does not assess implementation approaches in isolation. An assessor must consider the intersection of the three layers (the cyber security principle, the implementation approach and the protect surface) to determine what an effective implementation should look like.

The principle provides the intent. It describes **what** the enduring security outcome the organisation is trying to achieve. The implementation approach provides the action. It describes **how** that outcome should be realised in practice. The protect surface identifies the target. It defines the scope of **where** the approach must be applied and assess, including the specific system assets within it. The implementation level provides the sequencing. It defines **when** each protect surface should be addressed, directing organisations to apply principles and approaches to the most consequential surfaces first.

This three-way intersection is what makes the evolved framework different from a checklist. The same implementation approach can look materially different when applied to different protect surfaces, because the assets in each protect surface carry different risk profiles, operate in different contexts and face different malicious actor techniques. An assessor who evaluates an implementation approach without considering the protect surface it is being applied to may produce a rating that does not reflect the actual security posture.

When assessing a system, the assessor should ask three questions at each intersection:

- What does this cyber security principle require in the context of this protect surface?
- Does the implementation approach, as deployed, achieve that requirement for the assets within this protect surface?
- Is there evidence that the implementation approach is operating effectively and consistently across the protect surface?

Example: Vulnerability management

The vulnerability management cyber security principle includes an implementation approach requiring that patches or updates are applied within defined risk-based timeframes. When assessed against the 'identity services' protect surface, an assessor considers the specific assets in scope, which includes the domain controllers, federation servers, certificate authorities and MFA platforms. The assessor investigates whether the vulnerabilities in those assets are being identified and patched within timeframes proportionate to the systemic risk that identity infrastructure carries.

A 30-day patching cycle that may be acceptable for the 'business systems' protect surface would likely be insufficient for the 'boundary services' protect surface, where a single unpatched vulnerability on a web application server could allow initial access into an IT environment. The assessor rates the approach based on whether the patching program, as applied to web-facing infrastructure specifically, achieves the intended risk reduction and not whether the organisation has a patching program in general.

Example: Secure configuration management

The secure configuration management cyber security principle includes an implementation approach requiring that approved configuration baselines are defined, deployed and maintained. When assessed against the 'identity services' protect surface, an assessor considers the specific assets in scope, including directory services, federation platforms and authentication mechanisms. The assessor investigates whether approved configuration baselines have been defined that address identity specific risks, such as disabling legacy authentication protocols, enforcing LDAP bindings, restricting directory query permission and removing unnecessary services from domain controllers. The assessor then confirms that these baselines have been applied consistently across all identity infrastructure and that a mechanism exists to detect drift from the approved state.

When the same approach is assessed against the 'user devices' protect surface, the context changes. The assets in scope are corporate desktops and laptops, which is a larger and more diverse population. The assessor investigates whether approved configuration baselines have been defined that address endpoint specific risks and, critically, whether they are deployed consistently across the full fleet. An organisation may have well-defined baselines for Windows endpoints managed through Intune, but no equivalent for macOS devices. The assessor would rate

this as *partially effective*, as the implementation approach is in place for part of the protect surface but gaps in coverage limit its overall effectiveness.

Principle ratings within protect surfaces

Implementation approach ratings are aggregated to produce a single rating for each principle within each protect surface. The principle is not rated globally, but is rated separately for each protect surface to which it applies, recognising that a principle's effectiveness is determined by the assets and risk context of the surface in which it is being applied.

Within a single protect surface, a principle may have multiple applicable implementation approaches. The principle's rating at that surface is determined by the lowest applicable approach rating. If any approach under the principle at that surface is rated **not effective**, the principle is rated **not effective** at that surface. If all applicable approaches are rated as **effective**, the principle is rated **effective**. Where some approaches are **effective** and others are **partially effective**, the principle is rated as **partially effective**.

This rule preserves the integrity of the model at the level where it matters the most. The approaches under a principle work together to deliver a coherent security outcome at a given protect surface, and a gap in any approach undermines that outcome. The principle rating reflects the weakest link within the principle at that surface, not across the surface as a whole.

Example: Aggregating approach ratings to a principle rating

The identity, credential and access management principle has six applicable implementation approaches under the identity services protect surface. The organisation has rated four as **effective** and one as **partially effective**. The principle is rated **partially effective** at identity services. The organisation also rates the principle separately at boundary services, where all applicable approaches are **effective**, and the principle is therefore rated as **effective** at that surface.

Reporting posture

The framework provides a standardised format for reporting security posture to regulators, governing bodies and internal stakeholders. The reporting view is organised by implementation level and protect surface, expressing progress as the count of principles at each rating level.

This format demonstrates progression through the model without collapsing a protect surface to a single rating. An organisation that has improved several principles from **not effective** to **effective** can demonstrate that progress, even if some principles within the same surface remain **partially effective** or **not effective**. The format also avoids the misleading appearance of a single **not effective** rating disguising substantial work that has been completed across the surface.

The percentages shown in the below table are calculated using the following method: a Not Effective principle is worth 0, a Partially effective principle is worth 0.5, and Effective principle is

worth 1. the sum of those scores is then divided by the count of principles for a protect surface to get the reporting percentage.

In Visibility and Response, the organisation assesses:

- 5 principles as Effective (5 x 1 = **5**)
- 3 principles as Partially effective (3 x 0.5 = **1.5**)
- 1 principle as Not effective (1 x 0 = **0**)

The final calculation is the score divided by the number of principles

- 6.5 / 9 = **72%**

Table 2: Example report table of Organisation working towards Implementation Level 2

Level	Protect Surface	Effective	Partially effective	Not effective	% Principles Effective	Not assessed
All	Governance and culture	**5**	**0**	0	100%	0
All	Visibility and response	**5**	**3**	1	72%	0
1	Identity services	**8**	**4**	1	77%	0
1	Boundary services	**6**	**5**	0	77%	0
2	Critical process and services	**7**	**2**	0	89%	0
2	Privileged environments	**4**	**3**	2	61%	0
3	Business systems	**0**	**0**	0	0%	8
3	User devices	**0**	**0**	0	0%	9

The summary report should be accompanied by a detailed view that identifies which specific principles are rated **partially effective** or **not effective** at each protect surface, the implementation approaches that drove those ratings and the remediation activity in progress. The summary view demonstrates progression over time. The detail view exposes the specific gaps that require attention. Both views should be available to regulators and governing bodies to ensure that visible progress is not used to obscure outstanding risks.

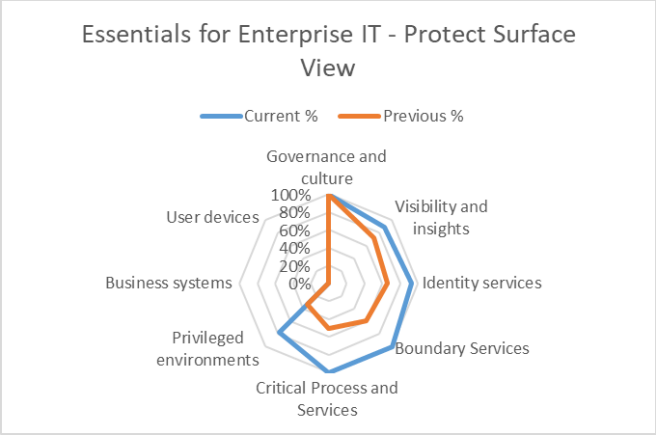
A principle that has not yet been assessed at a given protect surface should be recorded as **not assessed**. This is a neutral status that reflects sequencing, not failure, and should be distinguished from **not effective** in all reporting. Where a protect surface has not yet been addressed at all, all applicable principles will be deemed **not assessed** and the surface should be reported as such.

Visualising Posture and Progression

The reporting tables described above provide the structured detail required for assessment, regulatory submission and remediation planning. However, they are not always the best format for executive briefings, board reporting or rapid comparison across protect surfaces. The framework supports the use of complementary visualisations that present the same underlying data in a form better suited to those audiences.

A radar chart is particularly well suited to the measurement model because it expresses the multidimensional nature of an organisation's posture in a single image. Each axis of the graph represents the rating at that point. Overlaying successive assessments on the same graph produces an immediate visual representation of progression, making it straight forward to demonstrate where investment has moved the needle, and where further work is required.

Two views of the data are recommended for different audiences and purposes. A surface level view shows posture across all eight protect surfaces and is well suited to executive, board and regulator reporting. A principle-level view shows ratings across the principles applicable to a single protect surface and is well suited to working views for CISOs, security architects and assessors who are managing implementation at the surface level.



Visualisations should not be used in isolation. Aggregating ratings into a single point on a spider graph axis necessarily flattens the underlying detail, and a healthy-looking polygon can obscure individual gaps that warrant attention. Visualisations should always be presented alongside the structured tables that expose the specific principles and approaches driving the rating at each point.

Questions for you to consider:

- Do the proposed changes to the measurement framework seem logical and practical to

implement?

- What challenges or benefits do you see with the proposed measurement framework? If so, how could ASD support you to overcome these challenges?
- What additional reporting might be required for your stakeholders?
- How does the proposed assessment model align with your existing assessment processes?
- Are there particular protect surfaces or principles where you anticipate assessment ambiguity or differing interpretations?

Recognising existing Essential Eight investment

Organisations that have implemented the *Essential Eight maturity model will* have made significant and valuable investments in their security posture. The evolved *Essentials for enterprise IT* framework does not discard that work. The investment is carried forward; however, the mitigations and how we measure them have evolved and, in many cases, apply more broadly.

To support organisations to transition to the evolved *Essentials for enterprise IT* framework, ASD will provide additional guidance and support materials.

How the mitigation strategies map to cyber security principles

Each mitigation strategy under the *Essential Eight maturity model* can be mapped to at least one cyber security principle in the *Essentials for enterprise IT*. In most cases, a mitigation strategy maps to a cyber security principle that is broader in scope, meaning existing controls contributes to, but may not fully satisfy, the corresponding implementation approaches.

Application control maps to the 'trustworthy software and secure configuration management' cyber security principles. The evolved framework extends the concept beyond execution restriction to include integrity verification and control of administrative tooling, applied across multiple protect surfaces.

Patch applications and patch operating systems both map to the 'vulnerability management, asset identification and secure configuration management' cyber security principles. The evolved framework encompasses identification, prioritisation and remediation of vulnerabilities, and continues to advise on scenarios where patching is not feasible and compensating controls are required.

Restrict Microsoft Office macros maps to the 'trustworthy software and secure configuration management' cyber security principles. Specific macro controls remain valid but are treated as one element of broader approaches to restricting unauthorised code execution and maintaining approved configuration baselines.

User application hardening maps to the 'secure configuration management' cyber security principle. Organisations that have hardened web browsers, restricted web-based content, disabled unnecessary functionality and securely configured productivity applications have addressed controls that directly support the configuration baseline and drift detection implementation approaches.

Restrict administrative privileges maps to the 'identity, credential and access management, least privilege access, network segmentation and segregation and secure administration' cyber security principles. The evolved framework retains a strong focus on the separation of administrative environments from unprivileged environments, the use of time-bound privileges and the logging of administrative actions.

Multi-factor authentication maps to the 'identity, credential and access management' cyber security principle. The evolved framework continues to encourage the use of phishing-resistant

MFA and applies the requirement across multiple protect surfaces.

Regular backups maps to the 'regular and proven backups and least privilege' cyber security principles. The evolved framework continues an emphasis on immutable storage, isolation of backup repositories from production environments and regular restoration testing.

Event logging, detection and response capabilities that were required under the 'Multi-factor authentication**,** restrict administrative privileges, application control and user application hardening' map to several of the cyber security principles in the evolved framework, especially within the 'visibility and response' protect surface.

Assessing existing controls under the evolved framework

Organisations may find that a control which satisfied a requirement under the *Essential Eight maturity model* may require additional implementation to continue to be rated **effective** under the Essentials for enterprise IT. This is reflective of the evolving cyber threat environment and need for more sophisticated and targeted approaches to mitigate cyber threats. As the two frameworks measure different things, ASD recommends that existing controls be assessed through the evolved framework.

Breadth of cyber security principles

Each mitigation strategy from the *Essential Eight maturity model* may map to a cyber security principle, or multiple cyber security principles, that are broader than the original mitigation strategy and require additional work. This is new work, not lost investment. Existing controls will satisfy part of a cyber security principle, and the gap analysis process (detailed as Step 2 in the measurement section of this paper) will identify what remains to be addressed.

Outcome-based effectiveness

The *Essential Eight maturity mode*l historically accepted evidence of deployment, including whether a control was configured, a policy was applied or a setting was enabled. The evolved framework requires evidence that an implementation approach is achieving its intended effect. Organisations may need to supplement existing evidence with further operational evidence drawn from technical simulation testing and equivalent approaches, such as red team exercises, purple team activities, control validation testing, breach and attack simulation and penetration testing aligned to the specific implementation approaches being assessed.

Organisations currently assessed under the Essential Eight maturity model can use the gap analysis as the mechanism for translating their existing security posture into the Essentials for enterprise IT framework. Detailed mapping guidance will be provided in a future companion publication.

Questions for you to consider:

- Are you confident that your organisation could leverage the previous investments made under the Essential 8?

How the Essentials informs risk-based decisions

The framework is designed to mitigate the techniques used by malicious actors against Australian Government and CI operators. To support organisations in understanding why each implementation approach matters and where to direct their effort, most implementation approaches in the framework are aligned to one or more techniques from the MITRE ATT&CK knowledge base. The ATT&CK techniques represent the observed behaviours of real-world malicious actors and provide a common language for describing how compromises occur in practice. This alignment serves three purposes:

- Explains the security outcome each approach is intended to achieve in terms of specific malicious actor behaviours.
- Enables organisations to weigh implementation priorities against the threats most relevant to their sector and operating context.
- Allows organisations to integrate the framework with existing threat-informed defence practices, including red team exercises, threat hunting, and detection engineering.

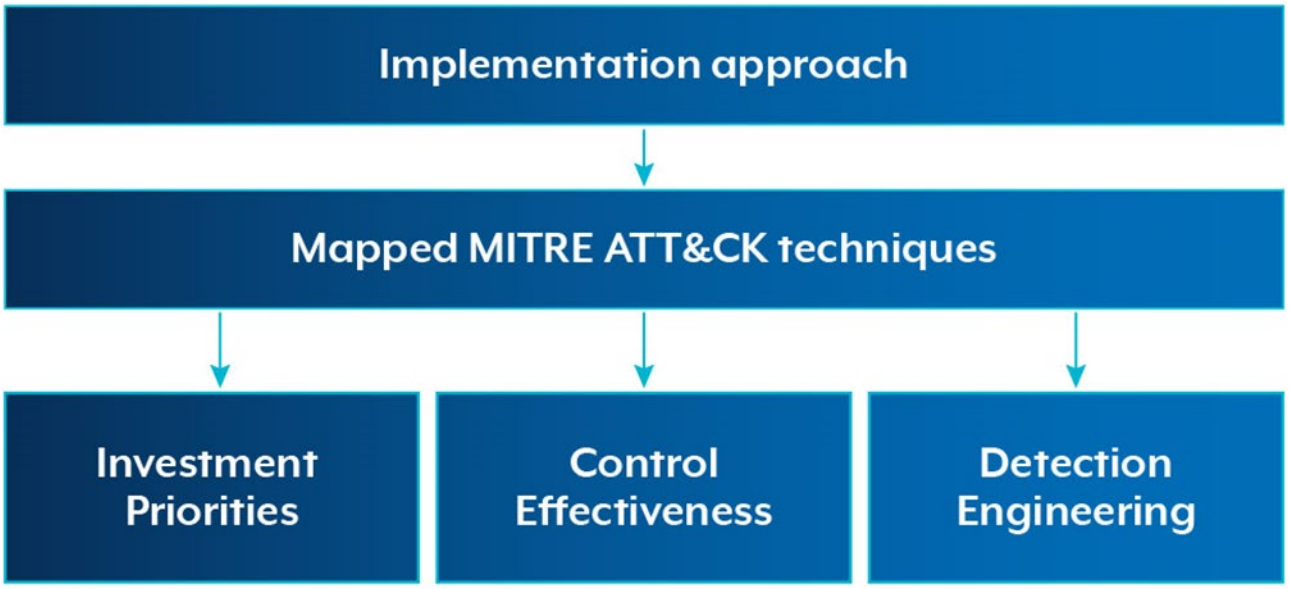


Figure 7: Essentials risk alignment purposes

How the mapping is structured

Most implementation approaches in the framework references one or more MITRE ATT&CK techniques that the implementation approach contributes to mitigating. The mapping is conservative, where an implementation approach is only mapped to MITRE ATT&CK technique where there is a direct relationship between the implementation approach and the disruption of that technique. Approaches that contribute peripherally to the mitigation of a technique are not listed against it. Some implementation approaches map to a single MITRE ATT&CK technique, while other approaches mapping to broad sets of techniques across multiple tactics.

Using the mapping to inform investment priorities

Organisations can use the mapping to weigh implementation priorities against cyber threat intelligence relevant to their sector. The recommended approach is as follows:

The organisation begins by identifying the threat groups, campaigns and techniques most relevant to its environment. This may draw on ASD threat advisories, sector-specific cyber threat intelligence, industry threat intelligence reporting or commercial threat intelligence subscriptions. The organisation then identifies the specific MITRE ATT&CK techniques most consistently observed in attacks against its sector or peer organisations.

The organisation then references the mapping to identify which implementation approaches contribute the most directly to mitigating those techniques. Where a small number of techniques account for a large share of the threats facing the organisation, the implementation approaches mapped to those techniques become the priority candidates for early investment. When developing a remediation plan for a **partially effective** rating, the ATT&CK mapping helps organisations determine whether a gap relates to commonly observed techniques or more sophisticated malicious actor behaviours. This informs the urgency of remediation.

This approach does not override the implementation level sequencing described earlier in this paper. IL1 surfaces should still be addressed before IL2 and 3 surfaces. The MITRE ATT&CK mapping informs prioritisation within an implementation level, not across implementation levels.

Using the mapping to validate effectiveness

The mapping also supports the validation of implementation effectiveness. Organisations with internal red team capability or that engage external testing services can use the mapping to design exercises that target the specific MITRE ATT&CK techniques an implementation approach is intended to mitigate. The result of these exercises provides operational evidence of effectiveness that complements the evidence developed during standard assessments.

Example: Trustworthy software

An organisation that had rated the *Trustworthy software* cyber security principle as **effective** under the ‘user devices’ protect surface can validate that assessment rating by attempting to execute a representative set of MITRE ATT&CK techniques mapped to that approach, including the execution of unsigned binaries and the use of unapproved interpreters. These activities were able to confirm that the control in place were successful in preventing and detecting the activity.

The mapping also supports detection engineering. Organisations operating mature security operations capabilities can use the mapping to identify gaps in detection coverage. Where an implementation approach is intended to mitigate a MITRE ATT&CK technique, but the organisation lacks detection logic that would identify attempted use of that technique, the gap should be addressed through detection rule development. This reinforces the relationship between the ‘visibility and response’ protect surface and the other protect surfaces.

Limitations of the mapping

The mapping is a tool to inform decision-making, not a substitute for it. Organisations should be aware of three limitations.

- **The MITRE ATT&CK knowledge base evolves continuously**. New techniques are added; existing techniques are revised and the relationships with mitigations change over time. The mapping for the implementation approaches will be updated periodically, but organisations relying on the mapping for high-stakes decisions should also reference the current MITRE ATT&CK knowledge base directly.
- **The mapping is necessarily generalised**. The effectiveness of any mitigation against any technique depends on how the control is implemented in the specific environment. An implementation approach mapped to a technique does not guarantee that the organisation's implementation of that approach will mitigate the technique in practice. The outcome-based measurement model detailed earlier in this paper remain the preferred basis for assessing whether a desired security outcome is being achieved.
- **Cyber threat intelligence is itself imperfect**. Malicious actors adapt their techniques, and prioritising mitigations based purely on observed historical activity risks underestimating the threats that are emerging or that has not yet been observed in the organisation's sector. The mapping should inform prioritisation alongside, not in place of, the broader risk management process.

Questions for you to consider:

- How would your organisation use the mapping to inform implementation?
- Are there any unintended consequences or expected benefits of the mapping guidance?
- What are the most useful ways that ASD can support your organisation to transition to the new framework?

Thank you for reviewing the consultation paper. ASD welcomes your comments and feedback on the proposed framework. The instructions for providing a written submission are available under the Purpose on page 2.

Glossary

Term	Meaning
baseline (behavioural)	A baseline of user activity, such as typical login patterns, resource access habits and anomalous actions.
baseline configuration	A documented configuration standard that has been formally reviewed and approved by an appropriate authority within the organisation.
boundary services (protect surface)	Systems positioned at network boundaries that broker connectivity between internal and external environments, including internet services and edge devices.
business systems (protect surface)	Supporting enterprise functions that collectively enable organisational operations.
compensating control	An alternative control that delivers the same intended security outcome as a primary implementation approach where that approach cannot be applied.
configuration drift	A change in the actual configuration of a system away from its approved baseline
critical function	A function performed by the organisation that is essential to the delivery of its primary outcomes or to the continued safe operations of its services
critical processes and services (protect surface)	Core processes and services that are essential to delivering primary organisational outcomes or contributing to personnel health and safety.
cross-cutting (technology)	A solution, capability or control that applies across multiple protect surfaces rather than being specific to one.
cyber security principle	An enduring cyber security objective drawn from the ASD information security manual.
cyber threat	Any circumstance or event with the potential to harm systems.
endpoint (technology)	A device used by personnel to access organisational systems and data.
federation (Identity)	A trust relationship between identity systems that allows authentication and authorisation in one system to be recognised by another

Term	Meaning
governance and culture (protect surface)	The policies, risk management processes, roles, responsibilities and security culture that underpin the cyber security program.
identity Services (protect surface)	Identity, credential and access management services including directories, authentication and authorisation enforcement solutions and credential stores.
implementation approach	A technology-agnostic description of a technical or procedural control that supports the achievement of a cyber security principle.
implementation level	A sequencing tier that groups protect surfaces to assist in prioritisation
implementation sequencing	The recommended order in which protect surfaces should be addressed when applying the framework
Indicators of Good Practice	A set of statements used to inform expert judgement about whether an implementation approach should be rated not effective, partially effective or effective.
internet-facing (technology)	A system, service or component that is reachable from the public internet, whether directly or through controlled ingress mechanisms.
lateral movement	The activity by which an adversary moves from a point of compromise to other systems within an environment.
legacy (technology)	Any technology that is no longer supported by the manufacturer, developer or vendor
MITRE ATT&CK	A public knowledge base of adversary tactics and techniques based on observed real-world behaviour, maintained by the not-for-profit organisation MITRE
principle-focus approach	A method of implementing the framework in which the organisation selects a cyber security principle and applies it across multiple protect surface
privileged environments (protect surface)	Systems and services that support privileged activities, management services and high-impact administrative functions.
protect surface	a logical grouping of enterprise functions and associated assets that share a common risk profile

Term	Meaning
residual risk (cyber)	The cyber security risk that remains after controls have been applied.
risk appetite	The level of cyber security risk an organisation is willing to accept in pursuit of its objectives
risk assessment (cyber)	The process of identifying cyber security risks, analysing their likelihood and consequence, and evaluating them against the organisation's risk appetite.
risk management (cyber)	The coordinated set of activities through which an organisation identifies, assesses, treats and monitors cyber security risks.
risk profile	A combined picture of cyber security risks an organisation is exposed to at a given point in time.
surface-focus approach	A method of implementing the framework in which the organisation selects a single protect surface and applies all applicable principles before moving to the next.
techniques (MITRE ATT&CK)	A representation of how an adversary achieves a tactical goal by performing an action
technology-agnostic	A description of what must be done without mandating a specific vendor, product or architecture
threat Intelligence (cyber)	Information about cyber threats, threat actors and adversary techniques that have been collected, analysed and contextualised to support decision making.
threat modelling	A structured activity to identify the threats relevant to a system, function, or organisation.
user devices (protect surface)	Endpoints issued to the workforce, including desktops and laptops.
visibility and response (protect surface)	The procedures and capabilities to detect, analyse and respond to cyber security events and incidents.

Annex A: Cyber security principles and implementation approaches

Summary

ASD is shifting to a broader, principles-based 'Essentials' series, evolving the existing Essential Eight to the Essentials for enterprise IT. The evolved framework moves away from requiring prescriptive technical controls, and instead provides a series of cyber security principles, with respective implementation approaches to address these principles

The appendix will assist executives and boards in their decision-making capacities regarding prioritisation, investment, and governance efforts. This appendix should be read in context of the overarching *Essentials for Enterprise IT* consultation paper that outlines the new approach in further detail.

Cyber security principles

In response to the evolving cyber threat and technology landscape, the framework builds on the original Essential Eight principles to deliver updated guidance aligned to contemporary tradecraft and modern enterprise IT environments. At a high level, it introduces:

- **Six new Govern principles** that establish accountability, resourcing, risk management and assurance over suppliers, artificial intelligence use and legacy systems. These principles ensure that the technical controls implemented across the framework operate within a governance structure that gives direction, visibility and longevity.
- **Two Identify principles** that address asset identification and security risk identification, ensuring organisations maintain an accurate inventory of systems, data and dependencies while continuously identify threats, vulnerabilities and business impacts. These principles are critical to an effective cyber operation as they establish a clear understanding of what must be protected and the risks it faces, enabling informed prioritisation, targeted controls and proactive risk management across the environment
- **Eight new Detect and Respond principles** addressing asset visibility, event logging integrity, detection logic, response procedures and post-incident learning. Together these principles articulate what a robust incident response capability looks like, highlighting the key role of security operations in modern environments, and acknowledge that detection and response must keep improving as threats evolve.
- **Twelve Protect principles**, eight of these adapted from the Essential Eight mitigation strategies and enhanced to cover modern adversary tradecraft and targeting. The remaining four principles are new additions that recognise tradecraft observed in recent compromises. Together the twelve protect function principles describe the technical hardening expected of organisations operating in the current threat environment.

The evolved framework is comprised of 28 cyber security principles with 58 associated Implementation Approaches, and assets and functions within Enterprise IT environments are grouped under eight protect surfaces.

How to use the matrix

The **Cyber security principles** outlined within the rows indicate the cyber security objective drawn from the ASD information security manual. These are broken down into **Implementation Approaches**, a technology-agnostic description of a technical or procedural actions that support the outcomes of a cyber security principle.

Where applicable, each principle maps to the malicious actor techniques (i.e. ATT&CK mapping) that it seeks to address.

The columns to the right of the ATT&CK mapping are the **protect surfaces**, the logical grouping of enterprise functions and associated assets that share a common risk profile. The matrix indicates which principles and implementation approaches apply to each protect surface. To address the implementation approaches as outlined in Table x, organisations will need to identify the relevant controls and apply those to all critical assets and functions, as scoped under the related protect surface.

The evolved framework introduces the concept of **implementation levels** (IL), a sequencing tier that groups protect surfaces to assist in prioritisation. Implementation levels define the recommended sequencing for addressing protect surfaces, prioritising systems that are most likely to be targeted first based on threat data.

- **Implementation level 1** (IL1) addresses boundary services' and identity services'. These two surfaces represent the entry point and the trust anchor of the environment.
- **Implementation level 2** (IL2) addresses 'critical processes and services' and 'privileged environments'. These two surfaces represent the next-highest-consequence assets in the environment and the elevated access pathways used to administer them.
- **Implementation level 3** (IL3) addresses 'business systems' and 'user devices'. These two surfaces represent the broader enterprise environment that supports day-to-day operations and the endpoints through which the workforce interacts with that environment.

Though the 'governance and culture' and 'visibility and response' protect surfaces are presented similarly to the other protect surfaces, these surfaces are applicable at every IL and the relevant principles should be considered in the context of the other protect surfaces.

As some cyber security principles are relevant to multiple protect surfaces, this table shows how the cyber security principles and implementation approaches correspond with the various protect surfaces. This table is useful to support the suggested gap analysis process and for prioritising the highest risk.

Cyber security principle	Implementation approach	ATT&CK mapping	Governance and culture	Visibility and response	Boundary services	Identity services	Critical processes and services	Privileged environments	Business systems	User devices
Executive cyber security accountability	Formally define and establish executive and board-level accountabilities for cyber security, including setting, monitoring and adjusting organisational risk appetite		✓							
Executive artificial intelligence accountability	Establish and maintain executive accountability for the use, governance and risk management of artificial intelligence within the organisation.		✓							
Cyber security resourcing	Assign sufficient cyber security personnel and resources with the required skills to operate, monitor, maintain and manage the lifecycle of controls across infrastructure, operating systems, applications and data.		✓							
	Periodically review cyber security resourcing against cyber threat trends, cyber security incident outcomes and organisational changes.		✓							
Security risk management assurance	Establish and maintain an approved risk management methodology and plan that is used to assess and treat cyber security risks		✓							
Supplier cyber security assurance	Establish and maintain a process for to only use software, hardware and services from trustworthy suppliers that have demonstrated a commitment to the security of, and transparency for, their products and services		✓							
	Establish and maintain a process for verifying that supplier cyber security practices are independently assessed on a regular basis		✓							
Legacy system management	Establish and maintain a process for identifying, assessing and managing the cyber security risk associated with legacy systems until they can be removed or replaced		✓							
Asset identification	Use an approved asset inventory mechanism that supports automated discovery to detect			✓						

Cyber security principle	Implementation approach	ATT&CK mapping	Governance and culture	Visibility and response	Boundary services	Identity services	Critical processes and services	Privileged environments	Business systems	User devices
	new and changed physical and non-physical assets.									
	Continuously identify and document infrastructure, operating systems, applications, identities, credentials and data as distinct assets.				✓	✓	✓	✓	✓	✓
Security risk identification	Identify relevant weaknesses, malicious actor techniques and cyber threat behaviours when undertaking risk assessments, and review these inputs on a continuous basis			✓						
Secure system lifecycle	Design systems to provide visibility, enforce authorised access, protect data, minimise required and default privileges, limit exposed services, and remove unnecessary functionality.	Valid Accounts Remote Services Exploitation for Privilege Escalation			✓	✓	✓	✓	✓	✓
	Secure build and deployment processes to enforce controls consistently across systems.				✓	✓	✓	✓	✓	✓
Cyber supply chain security	Verify the integrity and provenance of software, components and updates before they are introduced into systems.	Supply Chain Compromise Server Software Component Command and Scripting Interpreter			✓	✓	✓	✓	✓	✓
	Monitor systems for indicators of supply chain compromise and respond to identified concerns in a timely manner.	Supply Chain Compromise Server Software Component Exfiltration Over Web Service			✓	✓	✓	✓	✓	✓
Identity, credential and access management	Identify and disable or remove default and unused accounts on a regular basis.	Valid Accounts Brute Force			✓	✓	✓	✓	✓	✓
	Store authentication keys, secrets and tokens in approved secure repositories with access controls, including MFA	Unsecured Credentials Credentials from Password Stores				✓				
	Regularly scan codebases, configuration files, file shares and user collaboration platforms to identify unsecured keys, secrets and credentials	Valid Accounts Unsecured Credentials		✓						

Cyber security principle	Implementation approach	ATT&CK mapping	Governance and culture	Visibility and response	Boundary services	Identity services	Critical processes and services	Privileged environments	Business systems	User devices
	Use phishing-resistant MFA for users authenticating to systems, including from untrusted, external or high-risk locations	Create Account Valid Accounts Account Manipulation Brute Force Remote Services			✓	✓	✓	✓	✓	✓
	Disable legacy authentication protocols that bypass or weaken modern authentication controls	Brute Force Valid Accounts			✓	✓	✓	✓	✓	✓
	Ensure that user, machine, system and service credentials are unique per system or account, particularly in high-privilege contexts	Brute Force Valid Accounts			✓	✓	✓	✓	✓	✓
Least privilege access	Periodically review and update access privileges in alignment with current roles, duties and business requirements.	Account Manipulation			✓	✓	✓	✓	✓	
	Restrict permissions to create, modify, or delete accounts, roles, or access privileged functions or applications to authorised administrative roles.	Create Account Account Manipulation			✓	✓	✓	✓	✓	
	Provision user, administrative and non-person accounts with only the minimum access and privileges required to perform defined operational functions.	Create Account Account Manipulation			✓	✓	✓	✓	✓	
Secure administration	Ensure administrative activities can only be executed from dedicated and hardened environments that are isolated from unprivileged user workloads.	Exploitation for Privilege Escalation OS Credential Dumping			✓	✓	✓	✓		
	Enforce explicit session establishment controls before privileged actions are permitted, including the use of time-bound privilege grants.	OS Credential Dumping			✓	✓	✓	✓		
	Record all administrative or management actions that modify systems, configurations or controls with the details of the user, timestamp and target.	Impair Defenses Modify Registry Create Account Account Manipulation Create or Modify System Process			✓	✓	✓	✓	✓	
Secure configuration management	Define approved configuration baselines that disable or remove unnecessary services,	Server Software Component			✓	✓	✓	✓	✓	✓

Cyber security principle	Implementation approach	ATT&CK mapping	Governance and culture	Visibility and response	Boundary services	Identity services	Critical processes and services	Privileged environments	Business systems	User devices
	insecure settings, and unsupported or weak communication protocols.	Windows Management Instrumentation Network Service Discovery OS Credential Dumping								
	Deploy approved configuration baselines on systems, including infrastructure, operating systems and applications	Exploit Public-Facing Application External Remote Services User Execution Command and Scripting Interpreter			✓	✓	✓	✓	✓	✓
	Detect and respond to configuration drift from approved configuration baselines through automation or continuous technical assessments.	Impair Defenses Modify Registry			✓	✓	✓	✓	✓	✓
Vulnerability management	Identify known software and configuration weaknesses of assets through continuous and repeatable assessment activities, including vulnerability scanning activities	Exploit Public-Facing Application Exploitation for Privilege Escalation			✓	✓	✓	✓	✓	✓
	Apply security patches or updates within defined risk-based timeframes or, where operational impact is low, automatically.	Exploit Public-Facing Application Exploitation for Privilege Escalation			✓	✓	✓	✓	✓	✓
	Remediate or mitigate vulnerabilities exceeding defined risk thresholds within defined risk-based timeframes based on their severity and exposure.	Exploit Public-Facing Application Exploitation for Privilege Escalation			✓	✓	✓	✓	✓	✓
Trustworthy software	Only approved and integrity-verified software is permitted to execute on systems, with exceptions minimised risk-accepted, time-bound and regularly reviewed.	Impair Defenses User Execution Command and Scripting Interpreter Windows Management Instrumentation			✓	✓	✓	✓	✓	✓
	Restrict administrative and management tooling to approved software sets and executed from only authorised, known locations.	Impair Defenses User Execution			✓	✓	✓	✓	✓	✓

Cyber security principle	Implementation approach	ATT&CK mapping	Governance and culture	Visibility and response	Boundary services	Identity services	Critical processes and services	Privileged environments	Business systems	User devices
		Command and Scripting Interpreter Windows Management Instrumentation Masquerading								
Regular and proven backups	Backup important data applications and settings to immutable encrypted storage on a defined schedule to support timely restoration when required.	Data Encrypted for Impact			✓	✓	✓	✓	✓	
	Isolate backup repositories from general user environments and production workloads, ensuring they are separated from primary systems to prevent unauthorised modification or deletion.	Data Encrypted for Impact Data from Local System Data Staged			✓	✓	✓	✓	✓	
	Conduct regular testing of backup and restoration procedures to ensure that backups are reliable and capable of meeting defined recovery time objectives and recovery point objectives.	Data Encrypted for Impact			✓	✓	✓	✓		
Network segmentation and segregation	Isolate systems within protected network segments, restricting access from external networks and general user networks unless explicitly approved	Exploit Public-Facing Application External Remote Services			✓	✓	✓	✓		
	Limit lateral movement between, and within, protected network segments to explicitly approved pathways and services, based on business requirements and trust levels.	Remote Services Lateral Tool Transfer			✓	✓	✓	✓		
Content filtering	Inspect outbound content flows to detect unauthorised transfer of sensitive or high-value data.	Exfiltration Over C2 Channel Exfiltration Over Alternative Protocol Proxy			✓		✓	✓	✓	✓
	Block outbound network connections to known malicious or untrusted destinations.	Phishing User Execution Exfiltration Over C2 Channel			✓				✓	✓

Cyber security principle	Implementation approach	ATT&CK mapping	Governance and culture	Visibility and response	Boundary services	Identity services	Critical processes and services	Privileged environments	Business systems	User devices
		Exfiltration Over Alternative Protocol Proxy			✓				✓	✓
	Inspect inbound content from external or untrusted sources to prevent delivery of malicious or unsafe content.	Phishing Remote Services Ingress Tool Transfer			✓		✓		✓	
Cyber security awareness training	Ensure all personnel receive cyber security awareness training, that is tailored to how they will interact with systems, applications and data in their assigned roles, before being granted access to organisational systems and data.	Phishing Valid Accounts			✓	✓	✓	✓	✓	
	Update and provide cyber security awareness training on a regular basis to reflect current threats and operational considerations		✓							
Centralised event logging	Protect event logs from unauthorised access, modification or deletion by encrypting them during transit and at rest.				✓	✓	✓	✓	✓	✓
	Forward high-value, security-relevant event logs and configuration changes from infrastructure, operating systems and applications to a central logging capability.				✓	✓	✓	✓	✓	✓
	Establish a centralised event log storage capability that protects logs from unauthorised access or modification			✓						
Baselined high-risk access activities	Implement automated behavioural baselining to detect anomalous usage of legitimate software.	User Execution Command and Scripting Interpreter			✓	✓	✓	✓		
Cyber security event detection	Ensure event logs are normalised, time-synchronised and enriched to support correlation across systems and environments.				✓	✓	✓	✓	✓	✓
	Ensure detection logic identifies unauthorised changes, anomalous authentication activity and suspicious software execution.			✓						
Cyber security incident identification	Triage alerts and events in a timely manner to determine whether cyber security incident thresholds have been met.			✓						

Cyber security principle	Implementation approach	ATT&CK mapping	Governance and culture	Visibility and response	Boundary services	Identity services	Critical processes and services	Privileged environments	Business systems	User devices
	Use indicators from across identities, applications, data and infrastructure activity rather than isolated signals when performing cyber security incident identification activities.			✓						
Detection capability efficacy	Review detection coverage periodically to identify gaps arising from new assets, services or architectural changes.			✓						
Cyber security incident coordination	Integrate external coordination obligations into cyber security incident response activities where required.			✓						
Cyber security incident response	Use documented, tested and approved cyber security incident response procedures when responding to cyber security incidents.			✓						
Cyber security incident insights	Perform post-cyber security incident reviews to identify the security gaps that enabled each stage of the cyber security incident.			✓						
	Evaluate the monitoring coverage and cyber security incident response procedures following cyber security incidents.			✓						

Disclaimer

The material in this guide is of a general nature and should not be regarded as legal advice or relied on for assistance in any particular circumstance or emergency situation. In any important matter, you should seek appropriate independent professional advice in relation to your own circumstances.

The Commonwealth accepts no responsibility or liability for any damage, loss or expense incurred as a result of the reliance on information contained in this guide.

Copyright

© Commonwealth of Australia 2026

With the exception of the Coat of Arms and where otherwise stated, all material presented in this publication is provided under a Creative Commons Attribution 4.0 International license (https://creativecommons.org/licenses/by/4.0/).

For the avoidance of doubt, this means this license only applies to material as set out in this document.



The details of the relevant license conditions are available on the Creative Commons website as is the full legal code for the CC BY 4.0 license (https://creativecommons.org/licenses/by/4.0/legalcode.en).

Use of the Coat of Arms

The terms under which the Coat of Arms can be used are detailed on the Department of the Prime Minister and Cabinet website (https://www.pmc.gov.au/resources/commonwealth-coat-arms-information-and-guidelines).

Australian Government
Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE
ACSC Australian Cyber Security Centre