Foundations for modern defensible architecture

Australian Government

Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE

ACSC Australian Cyber Security Centre

Communications Security Establishment Canada

Canadian Centre for Cyber Security

Centre de la sécurité des télécommunications Canada

Centre canadien pour la cybersécurité

Te Tira Tiaki

Government Communications Security Bureau

National Cyber Security Centre

NEW ZEALAND

Bundesamt für Sicherheit in der Informationstechnik

国家サイバー統括室

National Cybersecurity Office

JPCERT CC®

警察庁

National Police Agency

National Cyber and Information Security Agency

NÚKIB

Table of contents

Introduction

Australian organisations, industries and individuals remain the target of malicious cyber actors. Cyber security continues to become more complex as organisations embrace flexible working, technologies rapidly develop, and the threat landscape evolves. The nature and persistence of threat actors targeting Australian networks require organisations to adopt a stance of 'when', not 'if', a cyber security incident will occur. The threat has made it increasingly difficult for network defenders to detect, prevent and respond to cyber security incidents. Organisations can take steps in the design, architecture and build of information environments to not only significantly minimise the risk and harm to their most critical assets and systems should an incident occur, but also to enhance their resilience.

This publication introduces modern defensible architecture (MDA) as a systematic and layered architectural approach to assist organisations in applying consistent, foundational aspects to build, maintain, update and enhance their systems.

The Australian Signals Directorate's Australian Cyber Security Centre (ASD's ACSC) and the following international partners provide the recommendations in this guide:

- Canadian Centre for Cyber Security – Canada
- National Cyber Security Centre – New Zealand
- Bundesamt für Sicherheit in der Informationstechnik – Germany
- National Cybersecurity Office – Japan
- JPCERT Coordination Centre – Japan
- National Police Agency – Japan
- National Intelligence Service – Republic Of Korea
- National Cyber and Information Security Agency – Czechia

Document purpose

This guidance has been developed to provide organisations with a baseline for secure design and architecture that prepare them to adapt to current and emerging cyber threats and challenges.

This guidance is drawn from ASD's experience in responding to cyber security incidents and performing security testing and vulnerability assessments of Australia's critical networks. It is based on considerations of technical cyber security practices such as zero trust and secure-by-design, which have emerged as better-practice approaches to increase cyber resilience.

Audience and scope

This publication is written for **technical security architects** and **enterprise architects** who are responsible for designing and building information technology (IT) environments.

While this publication provides recommendations targeted at securing corporate IT systems that support a workforce consisting of internal users, at a high level it is applicable to all types of environments.

This publication assumes an advanced level of computing and cyber security knowledge on the part of the reader.

Supporting materials:

Advice for **senior decision-makers**, outlining the benefits of MDA, can be accessed at the supporting publication. See *Modern defensible architecture for senior decision-makers*, at cyber.gov.au.

Advice for **managers** and **architects**, outlining how to develop a generational investment road map to implement modern defensible architecture, can be accessed at the supporting publication. See *Investing in modern defensible architecture* at cyber.gov.au.

What is modern defensible architecture?

Modern defensible architecture (MDA) is the name of ASD's ACSC's mission to ensure that organisations are considering and applying secure design and architecture in their cyber security strategy, resilience planning and implementations. It is based on the idea that there are certain elements of business and enterprise architecture that are common to any organisation that values cyber security.

MDA has been developed to assist organisations in preparing and planning for the adoption of technologies based on:

1. layered architecture and traceability as a methodical approach that separates security design into distinct levels, each addressing a specific aspect or scope of security management, from high-level business objectives down to specific technical implementations
2. zero trust principles of 'never trust, always verify', 'assume breach', and 'verify explicitly', implemented through zero trust architecture components and capabilities
3. secure-by-design practices that institute a security-first mindset within organisations when it comes to procuring or developing software products and services.

This document brings together these principles into 10 Foundations for modern defensible architecture (the MDA Foundations).

Many of the individual architectural MDA Foundations covered in this guidance are not new concepts but, when combined, they provide the ability to build a modern defensible architecture that is adaptable to emerging technologies and practices, and resilient to current and emerging cyber threats and challenges.

Organisations should regularly review their own architectural designs and decisions against each foundation to ensure they remain resilient and improve their cyber security maturity over time.

How do the MDA Foundations work?

Each Foundation represents an organisational security goal or capability that will facilitate a more efficient adoption of zero trust technologies and architecture. Organisations should create implementation roadmaps that prioritise the MDA Foundations that support their business objectives and protect their critical business capabilities and data. In doing so, organisations will improve their zero trust maturity and capabilities to achieve a modern defensible architecture. Implementing each Foundation contributes to a defence-in-depth approach to protect critical systems and data. This prevents or limits the impact of cyber incidents and the associated impact to critical business operations.

The MDA Foundations are designed to be technology agnostic. They allow organisations to make guided decisions on investment opportunities and design considerations, to identify technologies that are consistent with their requirements, and to take account of advancements in zero trust-enabling technologies and architecture.

The MDA Foundations recognise that every organisation is different, and the way they approach and prioritise implementation will be unique to individual organisational strategies and business objectives. The MDA Foundations are not in prioritised order, and organisations are encouraged to plan for the implementation of each Foundation as appropriate to their individual organisational context. For more details see *Investing in modern defensible architecture* at cyber.gov.au.

Designing and implementing architectural improvements to an enterprise environment will take significant time, resources and investment. Organisations should ensure that effort is applied to harden and protect existing systems by leveraging mature frameworks and prioritised mitigation strategies, such as ASD's *Essential Eight Maturity Model*. An organisation that works towards implementing a higher maturity level of ASD's Essential Eight will be well placed to adopt future guidance for achieving modern defensible architecture.

ASD has considered international advice and guidance on zero trust architecture alongside existing Australian government frameworks, including the *Protective Security Policy Framework* (PSPF); Hosting Certification Framework; cloud strategies; Gateway Standard and Guidance; and technical advice, including ASD's *Information Security Manual* (ISM) and *Secure-by-Design* publications.

Consultation feedback and key updates

The MDA Foundations were originally published as a consultation draft in February 2025. Feedback was received from more than 240 stakeholders from across government and critical infrastructure, ICT customers and vendors, and managed service providers. ASD thanks all stakeholders who participated in roundtable events and provided written feedback to better tailor this advice.

The MDA Foundations received positive feedback and its core message has not been altered. Updates have been made to better explain concepts, threat context and implementation actions. Key updates include:

New sections:

- a dedicated section to explain key underlying concepts.
- a single-page summary overview of all 10 MDA Foundations.

Updates to the 10 MDA Foundations:

- Some Foundations have been renamed to clarify their intent and scope.
  - Foundation 2 has been renamed from 'High assurance authentication' to 'High confidence authentication' in order to deconflict with ASD's other uses of the term 'high assurance'.
  - Foundation 8 has been renamed from 'Secure-by-Design software' to 'Secure-by-Design', to allow for a broader focus on a holistic, security-focused approach to business operations, functions and systems beyond just software.
  - Foundation 9 has been renamed from 'Comprehensive assurance and governance' to 'Comprehensive validation and assurance' to better focus on the assurance functions themselves, while allowing governance practices to sit outside the MDA Foundations.
- The MDA Foundations descriptions have been restructured for clarity.
  - The intent of each Foundation has been clarified with dedicated goal and overview sections.
  - New threat context overviews, referencing the MITRE ATT&CK framework, have been added with full threat explanations provided as an annex.
  - The Implementation Recommendations have been rewritten as Maturity Indicators, using control-based language to better target organisations' security objectives and integrate with other ASD publications.

If you would like to provide written feedback or have any questions regarding the MDA Foundations, please email us at acsc.sda@asd.gov.au.

Key concepts

Layered architecture and traceability

Layered architecture is a methodical approach that separates security design into distinct levels, each addressing a specific aspect or scope of security management, from high-level business objectives down to specific technical implementations. The top layers encapsulate strategic goals, capturing what the business aims to achieve and protect, while lower layers progressively define more precise and tangible security controls, technical implementations and operational measures. This layered approach ensures that security decisions are systematically aligned with organisational objectives, providing coherence and clarity at every stage.

Traceability refers to the clear linkage between each architectural level. It ensures that security controls and operational procedures directly support business goals and risk management strategies established at higher layers. This linkage allows for robust governance, whereby each security control can be justified by clearly traced origins in business objectives and threat context.

The MDA Foundations offer additional secure design and architecture advice as a structural framework upon which to implement ASD's ISM and ASD's Essential Eight Maturity Model. Properly implementing ISM controls and Essential Eight mitigation strategies remains important for mitigating targeted cyber intrusions and ransomware in IT environments. However, no set of mitigation strategies guarantees the prevention of all cyber security incidents, and both controls and mitigations are dependent on changes to technology and the threat environment.

Implementing mature security architecture will ensure information systems are able to maintain resilience over time and adapt as controls and mitigations evolve. This publication complements resiliency outcomes achieved by implementing mitigation strategies through controls and sound security architecture to increase networks resilience to cyber threats and challenges.

Figure 1 illustrates the relationship between ISM principles and strategic guidance, the MDA Foundations, and controls – practical guidance offered in both the ISM and Essential Eight. All layers are important to protect organisations from cyber threats and should be considered.

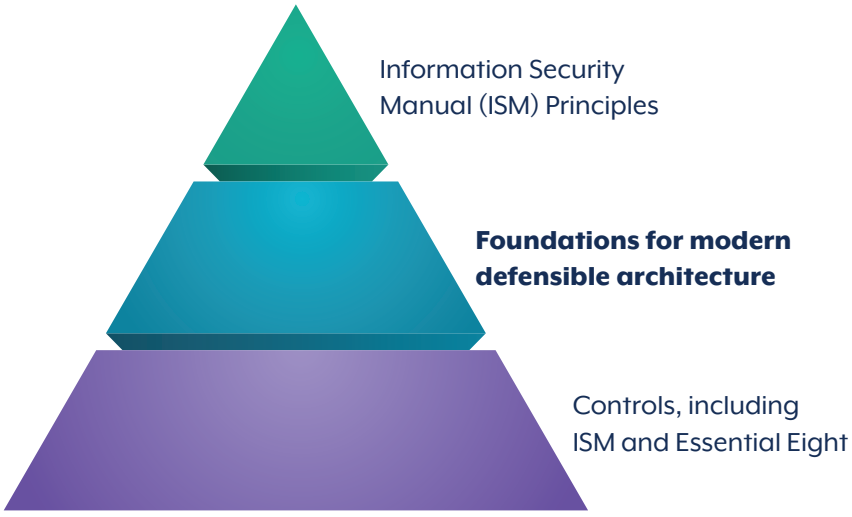

Figure 1. Layered architectural guidance

Zero trust and zero trust architecture

The following terms appear throughout the MDA Foundations, and are sourced from the US National Institute of Standards and Technology (NIST) – SP 800-207 – Zero Trust Architecture:

Zero trust provides a collection of concepts and ideas designed to minimise uncertainty in enforcing accurate, least-privilege, per-request access decisions in information systems and services in the face of a network viewed as compromised.

Zero trust architecture (ZTA) is an enterprise's cyber security plan that uses zero trust concepts and encompasses component relationships, workflow planning and access policies. Therefore, a zero trust enterprise is the network infrastructure (physical and virtual) and operational policies that are in place for an enterprise as a product of a zero trust architecture plan.

This publication refers to the US National Security Agency (NSA) **zero trust principles** (NSA – Embracing a Zero Trust Security Model):

1. **Never trust, always verify:** Treat every user, device, application/workload and data flow as untrusted. Authenticate and explicitly authorise each to the least privilege required using dynamic security policies.
2. **Assume breach:** Consciously operate and defend resources with the assumption that an adversary already has presence within the environment. Deny by default and heavily scrutinise all users, devices, data flows and requests for access. Log, inspect and continuously monitor all configuration changes, resource accesses and network traffic for suspicious activity.
3. **Verify explicitly:** Access to all resources should be conducted in a consistent and secure manner using multiple attributes (dynamic and static) to derive confidence levels for contextual access decisions to resources.

Authorisation model

A zero trust architecture treats every access request as potentially hostile until it can be justified. The authorisation model is the decision engine that judges each request against a real-time picture of business context and risk, then returns one of 3 outcomes: approve, deny, or approve with conditions, such as step-up authentication or session limits.

First, the model maps every enterprise resource (e.g. applications, data sets, APIs, devices) to the business functions it supports and the sensitivity of the information it holds. It then captures what 'legitimate use' looks like for each role, process, or machine account. This mapping establishes the baseline requirement that a subject – that is, the identity making the request – must meet to gain access.

When a user asks to open a resource, the model evaluates 4 context zones simultaneously:

1. **Identity context:** Who is asking? It checks the subject's role, employment status and clearance level.
2. **Task context:** Why is access needed right now and for how long? It looks at the workflow step the user is completing, the business transaction underway and any entitlements already granted.
3. **Environment context:** Where and how is the request originating? Factors include device health, network location, time of day, and any current threat signals or confidence signals (see next section).
4. **Data context:** What is the classification of the information being touched? Sensitivity, regulatory requirements and ownership determine allowable actions.

By grounding every decision in a dynamic business context, the authorisation model enforces least-privilege access without slowing legitimate work, and adapts continuously as environmental conditions change. Figure 2 shows an example of an authorisation model using business context to determine permitted access to resources.

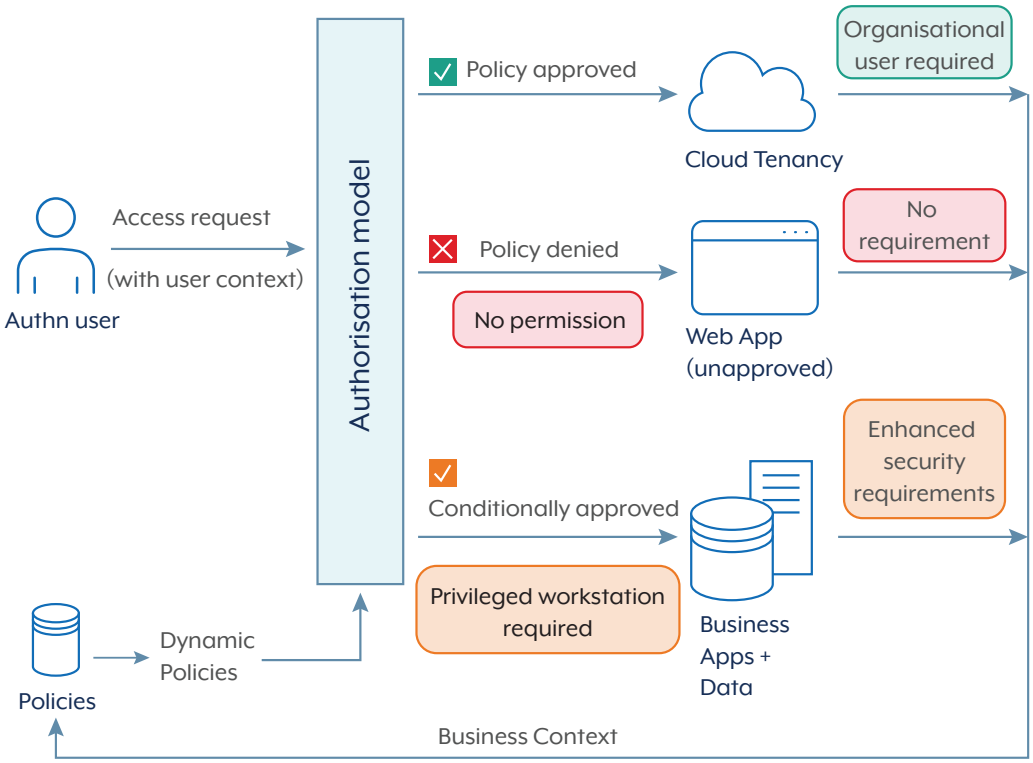

Figure 2. Authorisation model example

Confidence signals

A zero trust authorisation model moves beyond static credential authorisation by continuously evaluating access requests using dynamic confidence signals to grant or deny permissions. Confidence signals are near real-time indicators of trustworthiness and risk associated with the users and devices status and actions. These signals can be used to shape authorisation decisions.

Some key signals include:

1. **Device health:** Device compliance checks ensure devices meet defined security criteria, such as up-to-date endpoint detection and response software, secure configuration and encryption status. Healthy devices reduce the risk of compromise.
2. **Vulnerability scan data:** Regular vulnerability scans identify weaknesses in devices and software. Current scan data influences trust, with high-risk vulnerabilities triggering restricted access until remediation.
3. **Network location:** Network context, including IP reputation, geolocation and secure or unsecured networks, provides insight into potential risks. Access from known, expected or internal networks can increase the resultant confidence level, whereas unknown or suspicious locations should prompt caution.
4. **Behavioural analytics:** Analysing historical and current user activity, such as typical login patterns, resource access habits and anomalous actions, establishes behavioural baselines. Deviations from expected patterns signal potential security threats, prompting verification or restricted access.

By integrating these dynamic confidence signals, a zero trust model continuously assesses risks and enforces context-aware access controls, significantly enhancing security resilience and adaptive response capabilities to current and emerging threats. Figure 3 depicts an example of confidence signals integrating into an authorisation model.

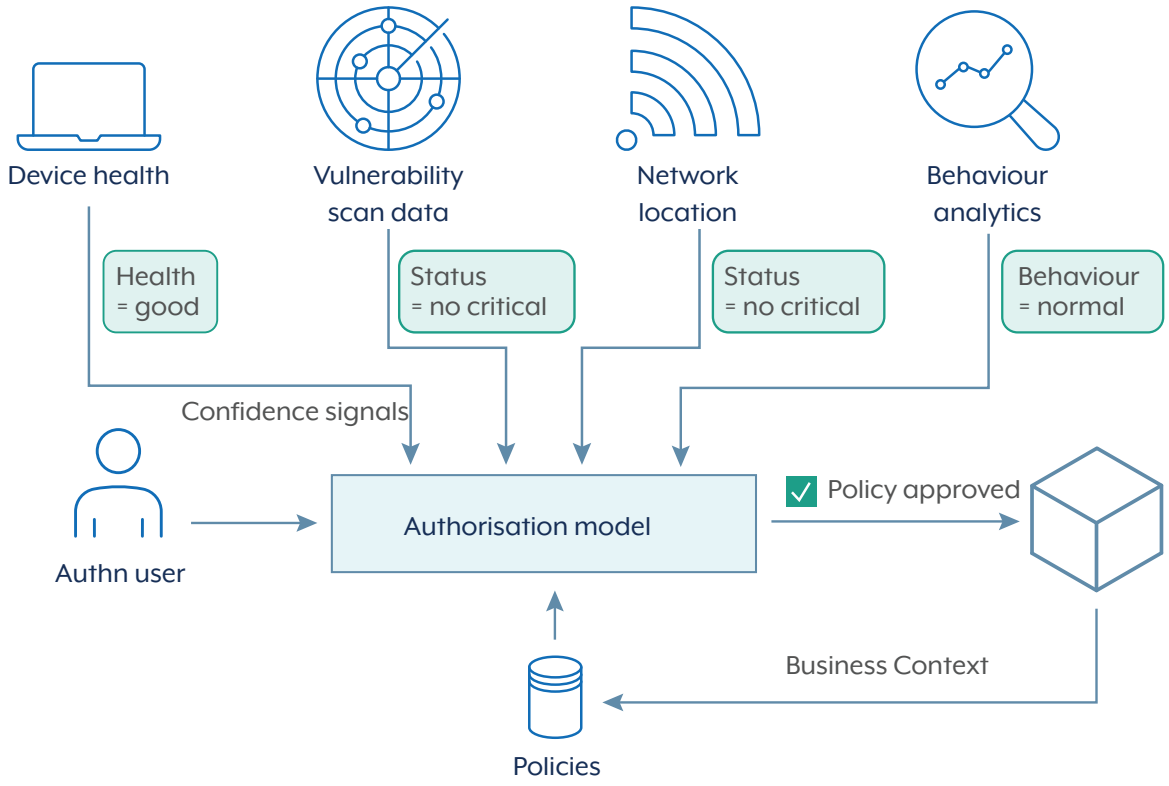

Figure 3. Confidence signals example

Policy enforcement points

In a zero trust architecture, trust is never assumed and should not be based on network location, device, or user identity alone. Every request is continuously evaluated, regardless of origin. Within this model, the policy enforcement point (PEP) acts as a gatekeeper that mediates access to resources.

PEPs may operate synchronously or asynchronously, depending on their function and placement within the architecture. In ZTA, PEPs are responsible for granting, denying, challenging or rejecting connections to resources. They support micro-segmentation solutions, which allow assets to be separated into their own security zones, ensuring that access is confined to the minimum necessary scope.

PEPs implement the decisions made by a policy engine, which is responsible for monitoring environment and behaviour in near real-time, and re-evaluate access decisions based on changing context. Factors relating to re-evaluation include device security posture, user behaviour, geolocation and threat intelligence feeds. Figure 4 shows an example decision flow undertaken by a PEP to enforce a security policy.

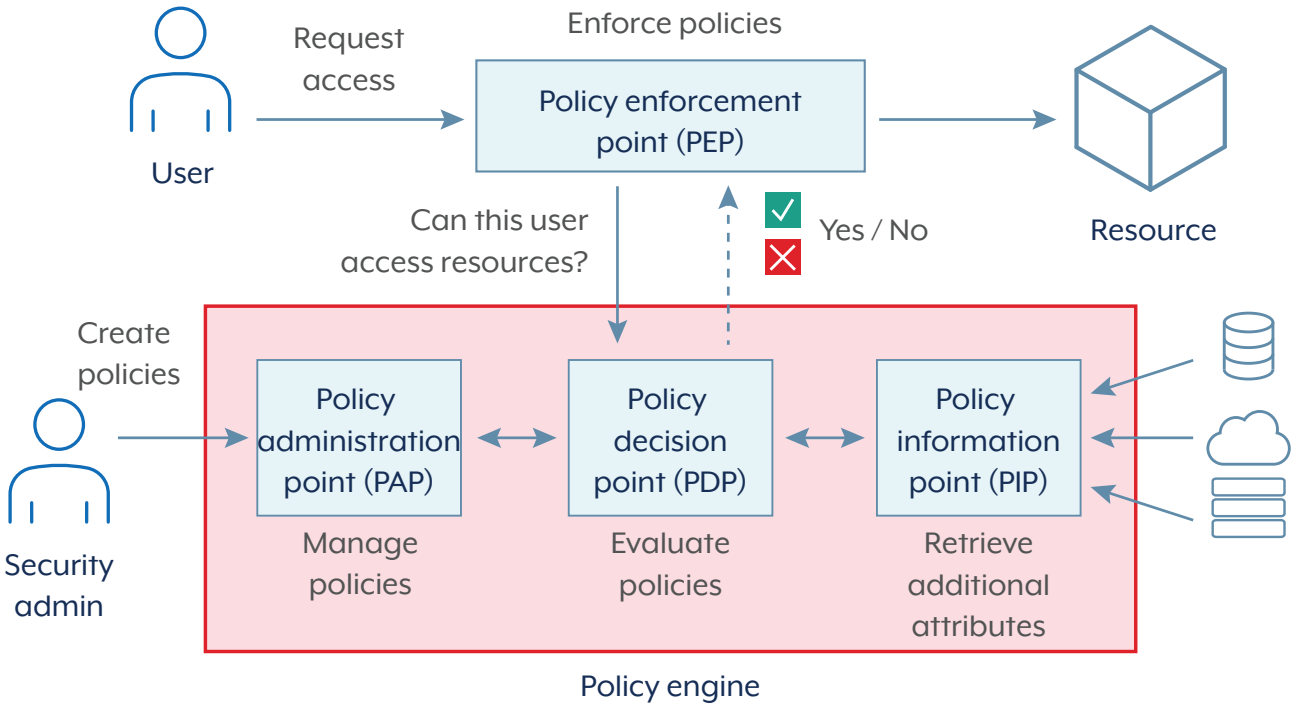

Figure 4. Policy enforcement points

The MDA Foundations

Summary of the MDA Foundations

Centrally managed enterprise identities

A reduced number of authoritative sources for enterprise identities by using centrally managed solutions. Enhanced visibility and accuracy of identities' registrations will reduce the likelihood and impact of identity compromise.

High confidence authentication

Strong and trustworthy authentication methods are used for all authentication events. Ensure that authentication events provide non-repudiation and prevent known authentication exploits. Re-authentication or access revocation is triggered when confidence in an entity drops.

Contextual authorisation

Authorisation to enterprise resources is initially and continuously validated, using the context of the sessions and resources to gain confidence in the request. High fidelity signals allow for higher confidence levels and more accurate authorisation.

Reliable asset inventory

A centralised repository has a complete and comprehensive knowledge of all endpoints, networks, applications, cryptographic assets and data stores that contain organisational information. Only with full visibility can an organisation truly assess the risks to the assets they own and operate.

Secure endpoints

Endpoints are resilient to cyber threats, compliant with organisational policy and continually provide contextual information to inform external systems. Trusting only validated resilient endpoints is key to limiting the impact of any compromise.

Reduced attack surface

A reduced number of possible attack surfaces that could be exploited. Limiting the avenues of attack allows for higher quality and concentration of mitigations in the correct places.

Resilient networks

Networks are tolerant to failure, resilient to attacks and protect data through the restriction of lateral (east/west) and vertical (north/south) movement to authorised requests. Protecting business functionality by limiting requests to only those authorised is vital for continued uninterrupted operations.

Secure-by-Design

Hardware and software are designed, built, validated, delivered and supported through security-first principles and practices, with known exploitable weaknesses reported and actioned throughout the life cycle of assets. Secure-by-Design's core value is to protect privacy and data and ensure security is maintained.

Comprehensive validation and assurance

Cohesion between business and security objectives is continually validated through assurance activities. Security measures must be validated to ensure their effectiveness as both a mitigation and as to support business functions.

Continuous and actionable monitoring

Real-time automated visibility and response using trusted high-fidelity and quality inputs. Without visibility organisations have less traceability and it will be harder to act when required.

Foundation 1: Centrally managed enterprise identities

Goal

A reduced number of authoritative sources for enterprise identities by using centrally managed solutions.

Overview

Identity management is foundational to secure operations and should be considered an organisation's primary security boundary control.

Identity defines how an entity is represented within the context of a digital system, typically represented by unique identifiers such as usernames, employee IDs or device serial numbers. In a zero trust environment, identity goes beyond user accounts to include devices, services and workloads as being represented by an identity that can be interpreted and managed. The accuracy and availability of identity management are foundational since every security decision relies on confidently knowing who or what is attempting to access a resource.

An organisation's identity management solutions should be compatible with sourcing identity data from external systems such as those used by personnel security, human resources and contract management teams. This data can include information about staff, contractors and third-party access to the environment. This approach provides organisations with a more holistic view of their user base, including their individual roles and responsibilities, which can assist in making organisation-wide risk-based decisions regarding user access requests and permissions. It also provides organisations with the ability to track users through organisational changes, including when staff change responsibilities or depart from their roles. Furthermore, chosen solutions should be compatible when sharing data with external systems, through methods such as authentication, authorisation and security monitoring solutions. These considerations should include organisational ways of working, such as cloud and hybrid solutions, to meet the organisation's modern business requirements and integrate with this context accordingly.

Threat context

Figure 5 shows the Mitre ATT&CK Tactics which have been identified as being strongly aligned to this Foundation, with further details found in Annex A of this document. Note that mapping is always context-specific. Organisations should refine this alignment further based on their own threat models and technology stacks.

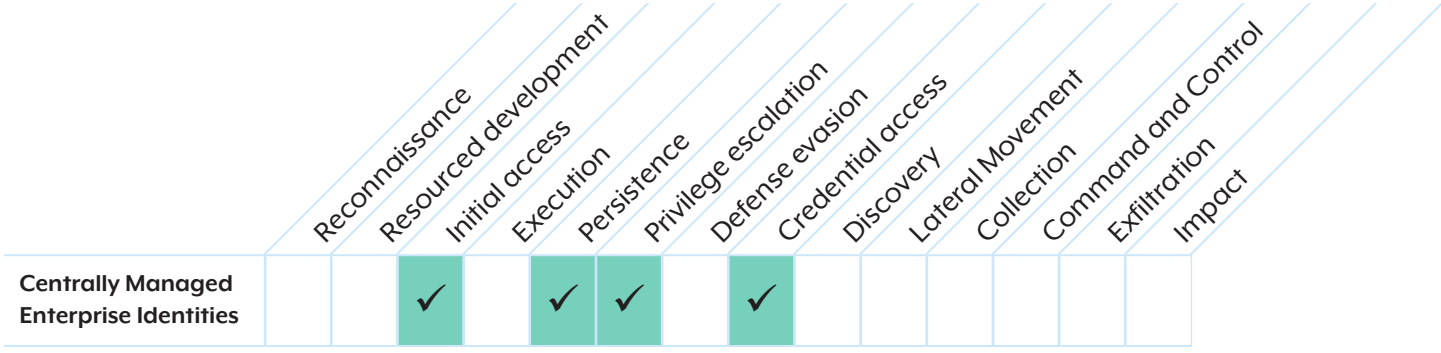

Figure 5. MITRE ATT&CK tactics identified as aligning with Foundation 1

Maturity indicators

The following maturity indicators have been developed for organisations to assess implementation maturity against the objectives of this Foundation. These indicators should be considered when making architectural design and investment decisions, alongside unique organisational risks, factors and mechanisms:

Identity management

- Enterprise identities are managed from the fewest number of sources possible.
- Actions taken on systems can be correlated with a single identity in the organisation's primary identity solution.

Least privilege

- All user and non-user identities are granted only the minimum privileges required to perform their functions correctly.

Separation of duties

- Users are prevented from performing a privileged action or opening a privileged session without approval from a secondary authoriser.

Federation and identity portability

- Identity solutions can provide identity data within and across security domains.
- Hardware, software and services are compatible with an organisation's identity management system.
- Identity solutions can consume data about identities from other security domains.

Life-cycle governance

- Non-user identities are captured and managed within centrally managed solutions.
- Identity data can be sourced from other systems within the environment.
- Identity information, both sourced and centrally managed, is accurately maintained.

Foundation 2: High confidence authentication

Goal

Strong and trustworthy authentication methods are used for all authentication events.

Overview

Authentication is the process of verifying a claimed identity and ascertaining that the entity requesting access is truly genuine. Traditional methods for authentication such as passphrases, SMS, email or mobile applications are replaced with stronger, phishing-resistant authentication methods, including passkeys or smartcards. Authentication for non-user identities should be achieved through factors that are tightly bound to the service or hardware, such as mutual authentication through known and validated digital certificates.

Cryptographic credential binding enhances the security of authenticated sessions for both managed and unmanaged endpoints, such as bring-your-own-device (BYOD). Organisations should prioritise investment in technologies that support cryptographic credential binding to both the device and the service being used.

Organisations need to consider the risk of compromise when choosing an authentication solution. This ensures that the chosen authentication solution can provide the necessary assurances that the claimant controls the authentication factors presented. Authentication requirements should be integrated up to the application layer, rather than just the network layer, to provide visibility and assurances of all authentication activities.

Threat context

Figure 6 shows the Mitre ATT&CK Tactics which have been identified as being strongly aligned to this Foundation, with further details found in Annex A of this document. Note that mapping is always context-specific. Organisations should refine this alignment further based on their own threat models and technology stacks.

Figure 6. MITRE ATT&CK tactics identified as aligning with Foundation 2

Maturity indicators

The following maturity indicators have been developed for organisations to assess implementation maturity against the objectives of this Foundation. These indicators should be considered when making architectural design and investment decisions, alongside unique organisational risks, factors and mechanisms:

Strong authentication

- Authentication solutions provide assurance that the claimant controls the authentication factors.
- Authentication is performed using cryptographically secure, phishing-resistant, multi-factor authentication methods.
- For non-user identities, authentication is performed using factors that are tightly bound to the service or hardware.
- Devices are required to authenticate to the network prior to being authorised to communicate with other resources.

Credential hygiene

- Authentication provides visibility and assurances of a user's particular activities.
- Communication between machines or services requires mutual authentication.
- Credential binding for authentication sessions is preferred.

Foundation 3: Contextual authorisation

Goal

Authorisation to enterprise resources is initially and continuously validated, using the context of the sessions and resources to gain confidence in the request.

Overview

Contextual authorisation considers each request to a resource in its full context including roles, permissions and contextual data. For each request, all available data points are evaluated to determine if access should be granted to a resource, which is done by an authorisation model. This is implemented using a combination of a PDP, PEP and PIP. For more information on these, see the earlier section on Key Concepts. In contrast, simple authorisation models only evaluate roles or permissions associated with the identity to determine access, and only on the first request in a session.

Contextual data can include attributes such as user's location, device security posture, or behavioural patterns. These data points, along with the principle of least privilege, need to be dynamic to respond to changing environmental factors, including new threat intelligence, device posture, or changes in observed user behaviour. As context data can change between requests, the confidence level will change and so may an authorisation decision.

Access must be continuously evaluated based on the activities taking place within a session, as well as other information that may indicate a drop in confidence, including changes in software state, security posture or behavioural analysis. Where real-time state data cannot be evaluated, historical state data that is used must be treated with a lower confidence. State data, including installed software configuration, that has drifted from recorded baselines should provide low confidence signals to the authorisation model, resulting in constraints or denials being applied to requests.

Organisations should standardise the level of assurance across their environment to the most secure option that suits the organisation's operational requirements. Additional local levels of assurance can be defined based on workload sensitivity and security requirements. For example, authorisation for some actions may be informed by users' mandatory training status.

Threat context

Figure 7 shows the Mitre ATT&CK Tactics which have been identified as being strongly aligned to this foundation, with further details found in Annex A of this document. Note that mapping is always context-specific. Organisations should refine this alignment further based on their own threat models and technology stacks.

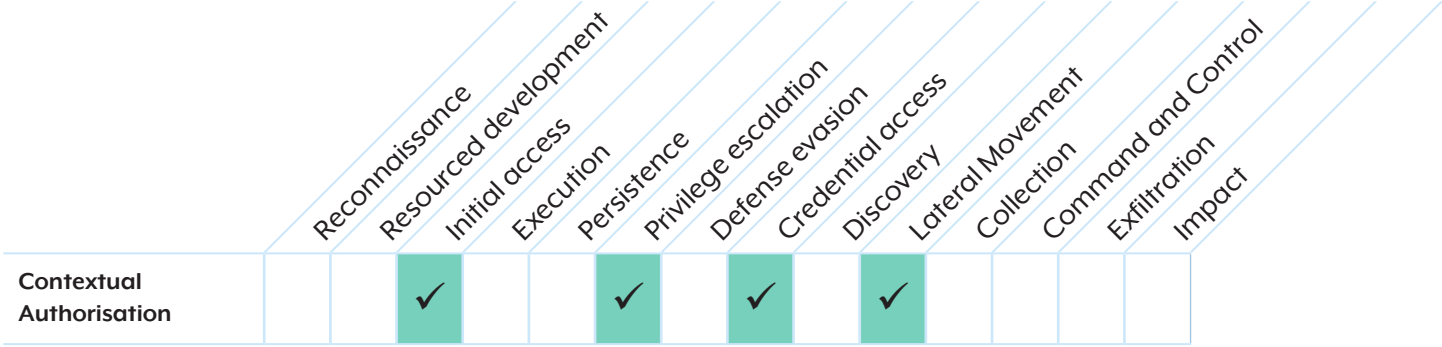

Figure 7. MITRE ATT&CK tactics identified as aligning with Foundation 3

Maturity indicators

The following maturity indicators have been developed for organisations to assess implementation maturity against the objectives of this Foundation. These indicators should be considered when making architectural design and investment decisions, alongside unique organisational risks, factors and mechanisms:

Authorisation model

- Authorisation solutions support a contextual authorisation model.
- Hardware, software and services provide contextual data to the authorisation model.

Secure authorisation policy development

- Dynamic authorisation policies adjust the level of confidence required for access based on broader organisational context.
- Authorisation policies include actions for when the contextual authorisation confidence level drops below the organisation's defined threshold.
- Authorisation policies are developed, stored and transferred securely to maintain integrity and availability.
- Logical system boundaries, where authorisation model decisions can be enforced, are defined.
- Policy enforcement points are located as close to resources as is practical.

Continuous access evaluation

- Authorisation for each session is granted individually, based on a minimum set of confidence signals to the required level of assurance, rather than inheriting trust directly from a previous session.
- The authorisation model compares contextual data from known devices against stored baselines.
- The expected values of session attributes are identified in access policies to provide indicators of confidence to the authorisation model.
- Authorisation is only provided for the minimal length of time required for that session.
- Established access is revoked if the confidence signals fall below that required for the session.
- Authorisation decision enforcement points generate and communicate ongoing environmental and session context for the continuous access evaluation.

Foundation 4: Reliable asset inventory

Goal

A centralised repository has a complete and comprehensive knowledge of all endpoints, networks, applications, cryptographic assets and data stores that contain organisational information.

Overview

Organisations need to have comprehensive knowledge of all their assets and their criticality. These can include endpoints, networks, applications and data stores that contain organisational information, including in services and environments not directly managed by the organisation. The asset inventory needs to be accurate and reliable, with the capability to continuously identify and record updated information on an organisation's assets and resources.

A comprehensive asset inventory will include details on all past and present organisational assets and information about the relationships between assets. The asset inventory will cover both physical and digital assets including endpoints, networks, systems, workloads, data and software. Dependencies of all assets need to have machine-readable traceability through software bill of materials (SBOM), hardware bill of materials (HBOM) and artificial intelligence bill of materials (AIBOM). Additionally, the asset inventory holds the requirements for authorised asset interactions and authorisations to access, transmit, store or process data. Decommissioned assets are marked as such in the asset inventory so that, along with other unapproved assets, they are not trusted when attempting to connect to the environment.

Automation is a critical component of an effective asset inventory and supports the discovery of new assets, changes to existing ones, and the removal of assets. Automation is a critical operational element in highly dynamic or short-lived environments, such as in systems that are designed to scale based on processing needs. Automation is not limited to monitoring and discovery but should include responsive actions such as blocking unapproved changes or assets and triggering alerts.

Threat context

Figure 8 shows the Mitre ATT&CK Tactics which have been identified as being strongly aligned to this Foundation, with further details found in Annex A of this document. Note that mapping is always context-specific. Organisations should refine this alignment further based on their own threat models and technology stacks.

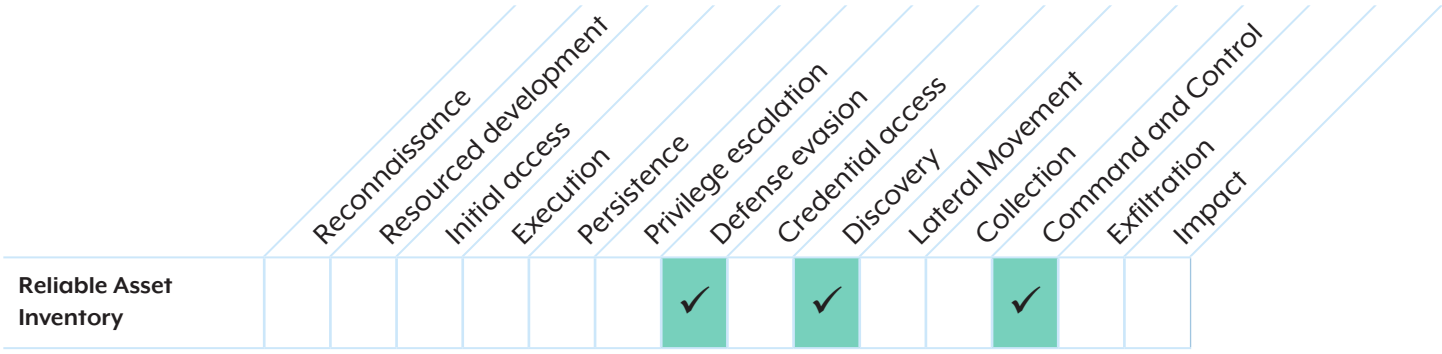

Figure 8. MITRE ATT&CK tactics identified as aligning with Foundation 4

Maturity indicators

The following maturity indicators have been developed for organisations to assess implementation maturity against the objectives of this Foundation. These indicators should be considered when making architectural design and investment decisions, alongside unique organisational risks, factors and mechanisms.

Inventory management

- Asset inventory information is stored in the fewest number of sources possible.
- Asset inventories store information on the criticality of assets.
- Asset inventories store information on the lifespan of assets.
- Asset inventories store information on the configuration state of assets.
- Asset inventories store information on the authorised interaction of assets.
- Asset inventories store information on the dependencies of assets.

Asset visibility

- The connection of new assets to the environment is automatically detected and reported to relevant teams and monitoring solutions.
- Asset inventories store and provide information in a machine-readable format.
- Asset inventory information is accessible to security software for analysis.

Asset lifecycle management

- Assets approaching end of life are reported to the relevant teams and monitoring solutions to ensure they are replaced.

For more information on SBOMs, please see A Shared Vision of Software Bill of Materials (SBOM) for Cybersecurity | Cyber.gov.au

Foundation 5: Secure endpoints

Goal

Endpoints are resilient to cyber threats, compliant with organisational policy and continually provide contextual information to inform external systems.

Overview

Organisations manage a variety of endpoints, including user devices, servers, networking equipment and virtualised environments. Each of these leverage a variety of hardware and software, both of which need to be protected and securely managed to prevent and minimise the impact of cyber threats. Before baselines are generated and recorded, endpoints that are not supported with secure-by-design 'loosening guides' should be hardened in line with relevant government and vendor hardening guidance. While this is often considered in the case of user endpoint devices, other endpoints such as servers and cloud infrastructure should also be secured in this way.

The operational effectiveness of endpoint baselines is highly dependent on their integrity and availability. Baselines for all classes of endpoints are required to allow organisations to continuously measure and monitor endpoints for unauthorised changes, and to respond appropriately to unauthorised changes. All endpoints need to be capable of generating and communicating quality and high-fidelity signals about security events and run-time configuration.

Threat context

Figure 9 shows the Mitre ATT&CK Tactics which have been identified as being strongly aligned to this Foundation, with further details found in Annex A of this document. Note that mapping is always context-specific. Organisations should refine this alignment further based on their own threat models and technology stacks.

Figure 9. MITRE ATT&CK tactics identified as aligning with Foundation 5

Maturity indicators

The following maturity indicators have been developed for organisations to assess implementation maturity against the objectives of this Foundation. These indicators should be considered when making architectural design and investment decisions, alongside unique organisational risks, factors and mechanisms.

System and application hardening

- Configuration baselines are built in adherence with the organisation's security objectives.

Configuration management and baselines

- Configuration baselines maintain their integrity through secure storage and communication.
- Authorised exemptions to baselines are tracked and regularly reviewed.
- Changes to endpoint configurations or operational settings are automatically detected and reported to relevant teams and monitoring solutions.
- Unauthorised deviations of endpoint configurations or settings from a recorded baseline are automatically restored.
- Endpoint context data is made available to the authorisation model.
- Endpoint configuration is made available to the monitoring solution.

For more information on ASD hardening guidance, please visit Guidelines for System Hardening | Cyber.gov.au.

Foundation 6: Reduced attack surface

Goal

A reduced number of possible attack surfaces that could be exploited.

Overview

By identifying and reducing the number of attack surfaces, there will be a reduction in the overall likelihood of successful exploitation by a malicious cyber actor. Additionally, reduced attack surfaces within endpoints, networks, applications and data stores can limit the impact of a breach by reducing the number of exposed surfaces to the initially compromised endpoint. Endpoints should be restricted from communicating between network segments unless required, as well as limiting the exposure of resources to networks where there is no requirement, or where a significant security risk is identified.

Organisations should consider the attack surface of systems they do not own or operate that contain organisational data, such as cloud and managed service providers, and externally hosted data repositories.

Attack surface reduction should be reflected in the organisation's asset inventory which can be used to enforce the restriction of connection assets. Internally managed applications that are required to be accessed remotely should do so in a way that is consistent with the organisation's authorisation model, and without relying solely on protections at the network layer.

Endpoints that have reached their end of life should be removed or replaced to further reduce organisational attack surfaces.

Threat context

Figure 10 shows the Mitre ATT&CK Tactics which have been identified as being strongly aligned to this Foundation, with further details found in Annex A of this document. Note that mapping is always context-specific. Organisations should refine this alignment further based on their own threat models and technology stacks.

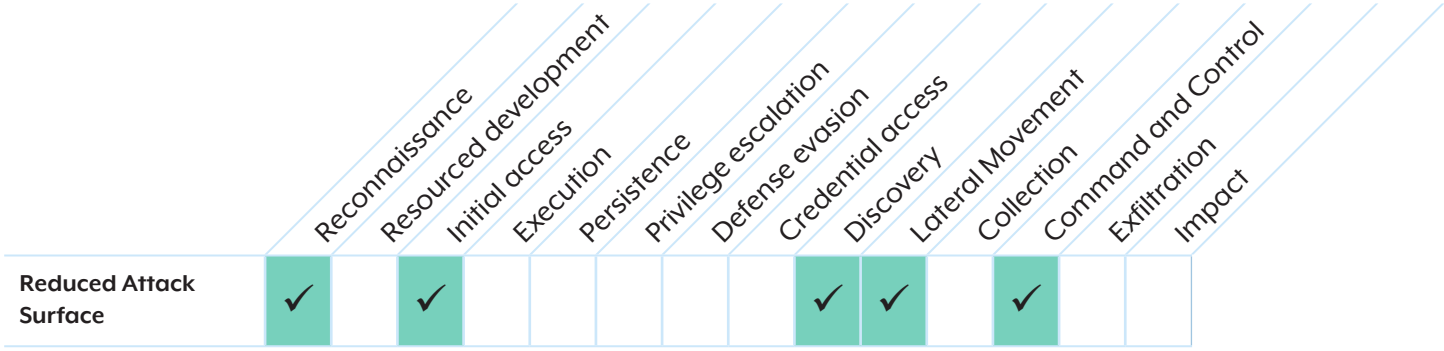

Figure 10. MITRE ATT&CK tactics identified as aligning with Foundation 6

Maturity indicators

The following maturity indicators have been developed for organisations to assess implementation maturity against the objectives of this Foundation. These indicators should be considered when making architectural design and investment decisions, alongside unique organisational risks, factors and mechanisms.

Service and port restriction

- External services that consume, process and store organisational data are identified and documented.
- Attempts to access unauthorised external services are automatically prevented and reported to relevant teams and monitoring solutions.
- Applications and networks are visible only to networks and resources that are required for them to perform their operations.
- The exposure of applications and networks to untrusted environments is monitored and automatically reported to relevant teams and monitoring solutions.

Software management

- Remote access to internally managed applications is consistent with the organisation's authorisation model.
- Endpoints have only the applications, features and functions installed or enabled to meet business requirements.
- A centralised and managed approach maintains the integrity of patches or updates and confirms that they have been applied successfully.
- Future depreciation and replacement of software and systems is planned for as part of procurement and monitored throughout its lifecycle.
- Vulnerable business-critical software which cannot be patched or upgraded is isolated from other resources when not being used, with the intent to replace the software as soon as possible.
- The deviation of software configuration or operational settings from the recorded baseline is monitored and automatically reported to relevant teams and monitoring solutions.
- The deviation of software configuration or operational settings from the recorded baseline is automatically restored to that baseline.

Vulnerability management

- The attack surface of operating systems and software is continuously monitored and mitigated.

Foundation 7: Resilient networks

Goal

Networks are tolerant to failure, resilient to attacks and protect data through the restriction of lateral (east/west) and vertical (north/south) movement to authorised requests.

Overview

Networks play a key role in both protecting organisational endpoints against cyber attacks and defending other resources when an endpoint is compromised. Network design and architecture provide security and support to organisational business requirements, noting both evolve organically over time and need to be regularly reviewed to maintain resilience.

Network mitigations are responsible for ensuring only authorised requests can transverse organisational networks, protecting the data in the requests from tampering, interception and exfiltration. Using the asset inventory, networks will define logical boundaries that will restrict which resources can connect to each other. This is done by using the context of the request to dynamically determine confidence in the validity of the connection. The network is also responsible for protecting services and systems availability and impact from events such as network failures, system compromises and denial-of-service attacks. These mitigations will assist in preventing compromised endpoints impacting other endpoints in the environment.

Chosen technologies need to support the organisation's operational requirements and user behaviours, including remote working and bring–your–own–device (BYOD), as part of risk-based decision-making that acknowledges how these may impact the security of software and data. All network technologies that are introduced into an organisation's environment should be designed with modern approaches to security. This includes the customisation of chosen protocols and ciphers, privacy, high availability and automated failover in the event of a device or service failure, and adverse risks of network technologies such as traffic capture and monitoring evaluated prior to their implementation.

Organisations can take proactive steps to reduce the business impact, improve network resiliency and identify architectural and operational flaws. Testing through table-top exercises, penetration tests, threat modelling, redundant network path failure testing, load testing and simulated device failure testing will support validation of design, architecture and implemented mitigations.

Threat context

Figure 11 shows the Mitre ATT&CK Tactics which have been identified as being strongly aligned to this Foundation, with further details found in Annex A of this document. Note that mapping is always context-specific. Organisations should refine this alignment further based on their own threat models and technology stacks.

Figure 11. MITRE ATT&CK tactics identified as aligning with Foundation 7

Maturity indicators

The following maturity indicators have been developed for organisations to assess implementation maturity against the objectives of this Foundation. These indicators should be considered when making architectural design and investment decisions, alongside unique organisational risks, factors and mechanisms.

Network segmentation

- Logical network boundaries are defined and monitored to detect and restrict lateral movement.
- Network device configurations are monitored, with unauthorised changes automatically reported to relevant teams and monitoring solutions.

Secure network design

- Network architecture is periodically reviewed to identify opportunities to improve network resiliency.
- Risks associated with the capture and decryption of network traffic are evaluated prior to use, and throughout the system's life cycle.
- Network technologies used are designed with modern approaches to security and privacy.
- Network technologies used are designed to securely support organisational business objectives.

Encryption and secure communications

- Organisational networks use proven secure and verifiable technologies.
- All data that is communicated over enterprise networks should be encrypted and configured to meet contemporary security standards and established better-practices, with weak or vulnerable network protocols deprecated as encrypted versions are standardised.

For more information on ASD-approved cryptographic protocols, please visit Guidelines for Cryptography | Cyber.gov.au.

Foundation 8: Secure-by-Design

Goal

Hardware and software are designed, built, validated, delivered and supported through security-first principles and practices, with known exploitable weaknesses reported and actioned throughout the life cycle of assets.

Overview

Secure-by-Design principles and practices are applied as a holistic security-focused approach to business operations and functions. Organisations need to be empowered by senior leadership to enable the whole organisation to adopt a security first approach. For the purposes of this Foundation, software should be understood to also include embedded software, firmware and their impact on hardware.

Software plays a critical part in an organisation's ability to meet business objectives and to stay secure against malicious cyber attacks. Software that is built with secure-by-design principles and practices is less likely to have exploitable weaknesses, which can reduce the likelihood of incidents occurring and the impact of any incident that does occur.

Software both internally developed and procured needs to be secure and verified before implementation. The supporting supply chain, including all software dependencies, needs to be validated and mitigations put in place to ensure that software is secure for its entire life cycle.

Software will be needed to support each of the MDA Foundations. Software needs to be compatible with the organisation's identity solutions, produce logs and telemetry for monitoring solutions, and be able to provide contextual data to the authorisation model.

Threat context

Figure 12 shows the Mitre ATT&CK Tactics which have been identified as being strongly aligned to this Foundation, with further details found in Annex A of this document. Note that mapping is always context-specific. Organisations should refine this alignment further based on their own threat models and technology stacks.

Figure 12. MITRE ATT&CK tactics identified as aligning with Foundation 8

Maturity indicators

The following maturity indicators have been developed for organisations to assess implementation maturity against the objectives of this Foundation. These indicators should be considered when making architectural design and investment decisions, alongside unique organisational risks, factors and mechanisms.

Compatibility

- Software is compatible with the organisation's identity solutions.
- Software is compatible with the organisation's authorisation model and can supply high fidelity contextual data.
- Software produces logs and telemetry compatible with the organisation's monitoring solutions.

Secure development and deployment

- The deployment of infrastructure and applications through software is immutable.
- The deployment of infrastructure and applications through software is idempotent.
- SecDevOps practices are embedded in organisational processes.
- Ongoing training and upskilling of development and operations teams is used to keep up to date with software advancements and prevalent cyber threats.

Supply chain security

- Procurement teams work closely with technical teams to determine the technical features and business requirements for new systems.
- Prior to the procurement of software, manufacturers and/or vendors are assessed for risks associated with their security practices, their reputation and their ability to quickly and effectively secure their products.
- Procured software is assessed prior to deployment to ensure it has been designed and built to be secure 'out of the box' against known prevalent cyber threats.
- Tooling used to analyse and implement technical policies is secured and verified before first use.
- Software dependencies are continually monitored throughout their life cycle for known vulnerabilities.

Threat informed decision-making

- Threat modelling is used in the development of all software.
- Threat intelligence is used to support threat modelling.

For more information on choosing secure and verifiable technologies, please visit Choosing Secure and Verifiable Technologies | Cyber.gov.au.

For more information on secure development, please visit Guidelines for software development | Cyber.gov.au.

Foundation 9: Comprehensive validation and assurance

Goal

Cohesion between business and security objectives is continually validated through assurance activities.

Overview

Organisations can only achieve cohesive security and business objectives when both are supported by high quality and consistent validation and assurance activities. Managing the relationship between security and business goals is an achievable goal when organisations recognise that they are not competing priorities, but a mutually connected pairing required for long-term operational success. Comprehensive validation and assurance activities that are reviewed and repeated regularly will help in achieving the goals of all MDA Foundations.

Validation and assurance activities don't just cover business policies and procedures, but include policies and procedures that integrate technical and security objectives. Authorisation processes need to account for both business and security objectives when considering the risks and benefits to the organisation. Definition and ownership of outstanding risks and validation of controls should be formally captured – including necessary technical governance and assurance testing – and feed into the organisation's accreditation and production-readiness assurance framework.

Threat context

Figure 13 shows the Mitre ATT&CK Tactics which have been identified as being strongly aligned to this Foundation, with further details found in Annex A of this document. Note that mapping is always context-specific. Organisations should refine this alignment further, based on their own threat models and technology stacks.

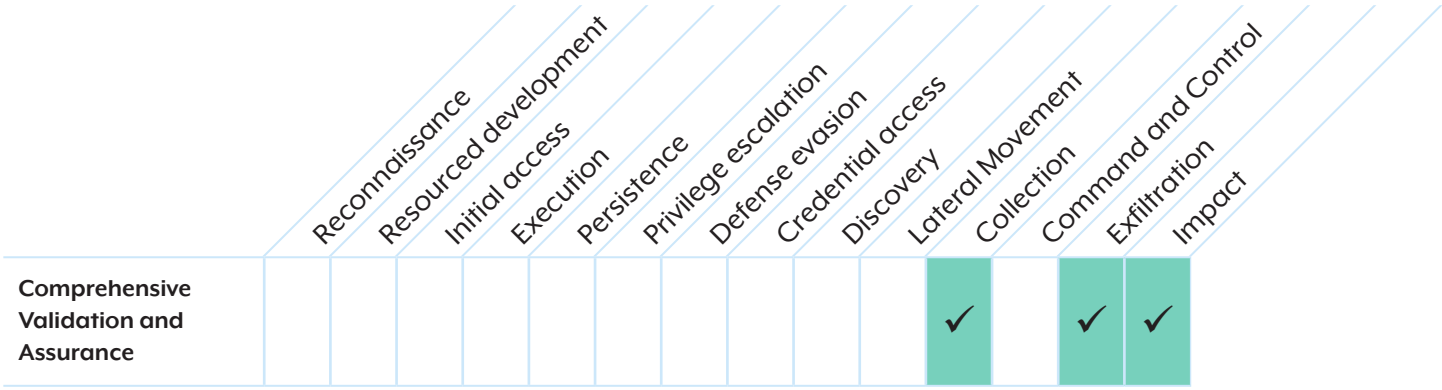

Figure 13. MITRE ATT&CK tactics identified as aligning with Foundation 9

Maturity indicators

The following maturity indicators have been developed for organisations to assess implementation maturity against the objectives of this Foundation. These indicators should be considered when making architectural design and investment decisions, alongside unique organisational risks, factors and mechanisms.

Security governance

- Assurance and governance activities are developed and delivered in accordance with relevant industry standards and government regulations.
- Assurance and governance activities outcomes are regularly reported to the appropriate level of senior management.

Assessments and continuous improvements

- The security and resilience of systems is assured and verified, both initially and on an ongoing basis, by appropriately skilled and equipped personnel.
- Assurance testing processes are repeatable and automatable.
- Assurance testing processes are run regularly with minimal human intervention.
- Outputs of assurance tests are stored securely and protected from compromise.
- Outputs of assurance tests are consumable by the authentication and authorisation models.
- Systems are regularly assessed for their resilience against current and emerging cyber threats and quickly detect changes that impact the security posture of the organisation.

Policy development and enforcement

- Technical policies for access authorisation and security baseline requirements are considered critical assets and are developed and stored on systems that are considered highly sensitive and hardened accordingly.
- Policies are regularly reviewed and updated based on cyber threat intelligence or changes in the environment.
- Policies are protected throughout their life cycle, with unauthorised changes, modifications or deletion actions automatically detected and reported to relevant teams and monitoring solutions.

Foundation 10: Continuous and actionable monitoring

Goal

Real-time automated visibility and response using trusted high-fidelity and quality inputs.

Overview

Visibility of all actions taking place within an organisation's environment or against organisational resources is required to identify any indicators of malicious activity. Monitoring solutions need to be capable of ingesting logs and contextual data from all systems and software within the organisation. The development of indicators of compromise needs to be completed with knowledge of organisational behaviours, risk tolerances and with trusted threat intelligence. These indicators can then be consumed to make informed or automated decisions on how to best respond. Fast response to indicators of compromise decreases recovery time and reduces the impact of a security incident.

Monitoring solutions should be capable of dynamic evaluation to support changes to the environment and to give weightings to the contextual data being consumed. The solution needs to be compatible with the organisation's asset inventory to allow for monitoring of all stored configurations against live configurations. This will ensure that security policies can be enforced effectively. Where configurations drift from registered baselines, alerts should be sent to the relevant teams. When drift is detected, automation should be set up to restore endpoints to baselines.

If indicators of compromise are detected, automated actions should be performed according to organisational risk tolerances. These could include removing, isolating, or restricting access to the compromised asset until the incident can be investigated and remediated.

Threat context

Figure 14 shows the Mitre ATT&CK Tactics which have been identified as being strongly aligned to this Foundation, with further details found in Annex A of this document. Note that mapping is always context-specific. Organisations should refine this alignment further based on their own threat models and technology stacks.

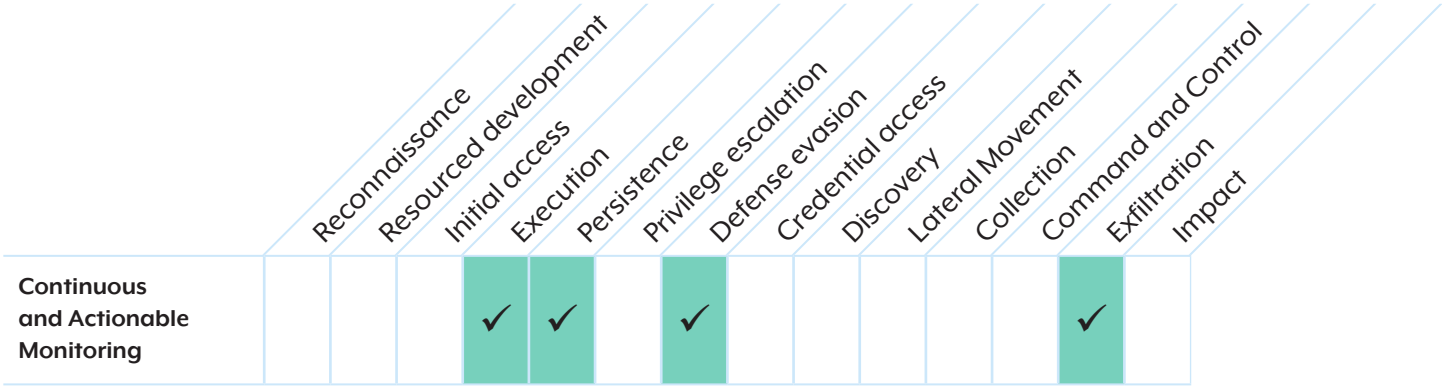

Figure 14. MITRE ATT&CK tactics identified as aligning with Foundation 10

Maturity indicators

The following maturity indicators have been developed for organisations to assess implementation maturity against the objectives of this Foundation. These indicators should be considered when making architectural design and investment decisions, alongside unique organisational risks, factors and mechanisms.

Centralised monitoring and logging

- Monitoring activities are informed by up-to-date and trusted threat intelligence sources, including from services that automatically deliver new data.
- All endpoints and software generate and communicate quality and high-fidelity signals about security events, including in logs, state and system telemetry, and behaviour.

Behavioural analysis and anomaly detection

- Monitoring systems are configured with organisational context to recognise behaviours and indicators that identify potential compromise, or likelihood that compromise could occur.
- Data sources used for alerting, monitoring and anomaly detection have a high assurance of integrity.
- Threat intelligence that includes indicators of compromise is analysed against existing systems as soon as is practical.
- Security monitoring solutions provide tailored alerts based on vendor or manufacturer advice regarding the patterns of behaviour in their software and services that could be an indicator of compromise.

Alerting and incident response

- The development and configuration of automated response actions are done in accordance with the organisation's risk appetite and tolerances.
- All actions triggered by automated response systems are monitored for further impacts and reverted to a known good state if unforeseen outcomes are detected.
- Systems and assets are automatically managed according to organisational risk tolerances when an indicator of compromise is detected.
- Response activities to incidents are automated to reduce the time taken to respond and the impact to the organisation.

For more information on monitoring and detection, please visit Guidelines for System Monitoring | Cyber.gov.au and Identifying and Mitigating Living Off the Land Techniques | Cyber.gov.au.

Annex A: Threat context

Overview

Figure 15 shows a summary of Mitre ATT&CK Tactics which have been identified as being strongly aligned to each Foundation. Note that mapping is always context-specific. Organisations should refine this alignment further based on their own threat models and technology stacks

Summary of MITRE ATT&CK tactics alignment

	Reconnaissance	Resourced development	Initial access	Execution	Persistence	Privilege escalation	Defense evasion	Credential access	Discovery	Lateral Movement	Collection	Command and Control	Exfiltration	Impact
Centrally Managed Enterprise Identities			✓		✓	✓		✓						
High Confidence Authentication			✓			✓		✓						
Contextual Authorisation			✓			✓		✓		✓				
Reliable Asset Inventory							✓		✓			✓		
Secure Endpoints				✓	✓			✓	✓					
Reduced Attack Surface	✓		✓						✓	✓		✓		
Resilient networks									✓	✓		✓		✓
Secure-by-Design		✓		✓		✓	✓							
Comprehensive Validation and Assurance											✓		✓	✓
Continuous and Actionable Monitoring				✓	✓		✓						✓	

Figure 15. Summary of MITRE ATT&CK tactics alignment

Alignment by Foundation

Foundation 1: Centrally managed enterprise identities

Mitre ATT&CK Tactic	Alignment to Foundation 1: Centrally managed enterprise identities
Initial Access (TA0001)	Adversaries who gain initial access to an organisational security domain through compromised identities often gain access through accounts externally managed by a trusted partner, such as a third-party vendor or service provider, that are insecure. Organisations that have centralised the management of organisational identities will increase the difficulty and noise of an adversary attempting to achieve initial access. In addition, it will reduce the number of identities that can be targeted and used for further techniques and access.
Persistence (TA0003)	Adversaries seeking persistence may add or modify local, domain or cloud accounts, use long-lived secrets and cycle passwords to avoid detection and credential resets. Organisations that centralise the creation and configuration of identities will have greater visibility, be more resilient and have a quicker response to unauthorised changes.
Privilege Escalation (TA0004)	Adversaries often require significant user privileges within a security domain to enable further techniques. They are likely to target identities with high privileges or attempt to raise the privilege of an already compromised account. Organisations who centrally manage the provision of privileges and roles will be able to reduce the likelihood of privilege misuse or privilege escalation, especially through third-party managed privileged identities.
Credential Access (TA0006)	Adversaries pursuing access can harvest or guess credentials, hashes or tokens through a variety of techniques to impersonate legitimate users and services. Centralised identity management gives organisations a single authoritative directory covering all users, devices and services with enforced MFA, least-privilege roles and shared telemetry. Any breach allows security teams to revoke account access rapidly and anomalous logins become obvious. This raises the cost of brute-force, phishing and token-theft techniques, while giving defenders a single vantage point to detect unusual password resets, token requests or directory changes.

Foundation 2: High confidence authentication

MITRE ATT&CK Tactic	Alignment to Foundation 2: High confidence authentication
Initial Access (TA0001)	Adversaries can use valid accounts to gain an initial foothold in an environment. These accounts can be compromised through several means, including phishing, default credentials, exploiting public-facing applications and poor secrets management. Organisations that work towards all accounts using secure and high confidence authentication such as MFA and credential binding will be more resilient to adversaries gaining access to valid accounts.
Privilege Escalation (TA0004)	Adversaries seeking privilege escalation can exploit poorly configured roles, steal tokens or exploit vulnerable authentication services to leap from a low-level account to administrator rights. This gives adversaries permissions to devices and platforms that have access to sensitive data. Organisations that use secure authentication for privileged identities (including non-user identities) will be better protected against common attacks.
Credential Access (TA0006)	Adversaries will attempt to compromise legitimate user credentials to gain access to organisational services and resources. Adversaries have several techniques for compromising existing credentials, including phishing, replay attacks, brute-force guessing, key logging and memory dumps. Organisations that use phishing-resistant and cryptographically device- bound authentication mechanisms will significantly reduce the risk of credential access exploitation.

Foundation 3: Contextual authorisation

MITRE ATT&CK Tactic	Alignment to Foundation 3: Contextual authorisation
Initial Access (TA0001)	Adversaries will commonly achieve initial access by using valid accounts with compromised credentials, or by exploiting public-facing applications. Organisations who have implemented a strong contextual authorisation model will be able to use contextual data to create a confidence level that is used to either grant or deny access to a requested resource.
Privilege Escalation (TA0004)	Adversaries may try to gain increased privileges and permission using compromised identities. An adversary can make attempts to manipulate existing permissions, circumvent existing control mechanisms or use system exploitation. Organisations that use behavioural analysis within a contextual authorisation model can better prevent identities from performing activities that fall outside expected behaviours. This includes any actions that are not typically performed by unprivileged identities.
Credential Access (TA0006)	Adversaries can use common techniques, like credential brute force attacks or input capture to gain access to the credentials of valid accounts. These techniques can leave indicators, which are picked up by logs & monitoring solutions, such as Security Information and Event Management (SIEM) platforms. An organisation can use these solutions and platforms to feed confidence signals to a contextual authorisation model, which can temporarily disable or prevent resource access from identities that are suspected of being compromised.
Lateral Movement (TA0008)	Adversaries will try to pivot to other endpoints in the environment to gain better visibility, achieve their objectives or find a place to hide and maintain persistence. Adversaries will use various techniques to move between connected systems, including further software exploitation or using legitimate tools and applications, which is often made simpler if considered trusted on the network. An organisation can reduce the opportunity for adversaries to move between systems by removing the inherited trust between network endpoints. Organisations who use the network source and destination to define acceptable attributes within the contextual authorisation policies will have finer control over remote endpoint access requests.

Foundation 4: Reliable asset inventory

MITRE ATT&CK Tactic	Alignment to Foundation 4: Reliable asset inventory
Defense Evasion (TA0005)	Adversaries will try to avoid detection by organisational security tools and platforms by using techniques that obfuscate their activities or their location within the environment. An adversary can create new endpoints – such as virtual machines, containers or applications – in the environment to avoid security policies and detection tools. Organisations with a comprehensive asset scanning capability will have a better chance of detecting any unauthorised creation of new computing resources.
Discovery (TA0007)	Once inside a network, adversaries probe for unmanaged hosts, shadow IT and forgotten services that might yield further access. A continuously reconciled asset inventory exposes any device or workload that appears without a record, feeds alerts to security tooling and lets defenders quarantine 'unknown unknowns' before the attacker can exploit them. By ensuring that every asset is accounted for, organisations deprive attackers of blind spots to pivot toward.
Command and Control (TA0011)	Adversaries commonly use infrastructure under their control to stage and launch attacks on the organisation's environment. This infrastructure can belong to the adversary or be the outcome of a previous compromise. Organisations that have implemented a complete and reliable asset inventory are able to map assets against the details of trusted services and endpoints that they are allowed access to, enabling blocking of unauthorised connections.

Foundation 5: Secure endpoints

MITRE ATT&CK Tactic	Alignment to Foundation 5: Secure endpoints
Execution (TA0002)	An adversary will often run malicious code on endpoints that have been already compromised or have vulnerable software installed. An adversary will aim to run malicious code on a host, either by tricking users, abusing in-built interpreters or starting services allowing them to further their objectives. An organisation that has hardened its endpoints and applications will be able to better detect and prevent the execution of malicious code.
Persistence (TA0003)	An adversary who has gained initial access to an endpoint will often try to maintain persistence by modifying the system to allow access, despite any interruptions or attempts to remove them from the environment. An adversary may change the security configuration of the operating system, register keys or replace drivers with malicious versions to create exploitable weaknesses that can give them another way to gain access or modify the system behaviour that reduces the likelihood that their activities will be detected. An organisation that can detect changes in the baseline configuration of a system and revert it to its former state will make it more difficult for an adversary to maintain persistence through these methods.
Credential Access (TA0006)	Endpoints are prime hunting grounds for cached passwords, browser cookies and authentication tokens. Adversaries have several techniques for compromising credentials, including key logging and memory dumps. Organisations can implement hardening measures such as memory-protection features, credential-guard technologies and encrypted storage to restrict an adversary's ability to dump credential material or scrape secrets from process memory. In parallel, endpoint baseline-drift monitoring can highlight any security control that has been disabled in pursuit of credential theft, enabling swift remediation.
Discovery (TA0007)	Once an adversary has compromised an endpoint through exploitation of a vulnerability, they will typically undertake discovery techniques to identify areas of opportunity, or to find a path to their objective. To reduce the chance of detection and enable their activities, an adversary will often try to modify the security configuration of the endpoint. An organisation that uses endpoint health as a confidence signal into an authorisation model will be able to restrict endpoint access to those who have reported misalignment with approved baselines, or endpoints that do not have an updated health status

Foundation 6: Reduced attack surface

MITRE ATT&CK Tactic	Alignment to Foundation 6: Reduced attack surface
Reconnaissance (TA0043)	Adversaries performing reconnaissance gather information on exposed services, legacy protocols and misconfigured systems to plan their attacks. The more information an adversary has on a target network, the easier it is for them to plan their path to exploitation. It increases the probability that they will find an unmonitored asset and allow intrusion campaigns to be precisely tailored, greatly amplifying organisational impact. When an organisation deliberately reduces its attack surface through the elimination of unnecessary public-facing services, unused services, restriction of endpoints to business-critical applications, patching of exposed software, limiting open ports, and ensuring networks are not exposed to lower-trust environments, the information available to an adversary during reconnaissance is reduced. This lowers the likelihood and accuracy of subsequent targeting attempts.
Initial Access (TA0001)	Adversaries seeking to gain initial access will probe internet-facing hosts, spearphish users, or exploit vulnerable or misconfigured open ports to gain their first foothold. Each surplus interface or legacy protocol multiplies organisational exposure, providing attackers with numerous parallel entry lanes and reducing defenders' ability to patch or monitor every ingress point before compromise occurs. An organisation that properly reduces attack surfaces through disabling default interfaces, enforcing strict boundary controls, and adopting contextual inbound filtering will minimise the number of externally reachable vectors, making successful initial access far less probable.
Discovery (TA0007)	Adversaries will sweep networks and systems to catalogue hosts, services, shares and trust relationships to map pathways for lateral movement or privilege escalation toward sensitive assets. Situational awareness from discovery techniques lets attackers avoid well-defended zones, focus on soft targets and chain weaknesses together, raising the probability of undetected escalation and intensifying the eventual business impact. Organisations that apply host-based firewalls, remove legacy administrative shares, tightly scope service exposure, and reduce attack surfaces limit what can be discovered internally, delaying attacker progression and increasing detection opportunities.

MITRE ATT&CK Tactic	Alignment to Foundation 6: Reduced attack surface
Lateral Movement (TA0008)	Adversaries seeking to pivot from their first foothold to new hosts rely on reusing stolen credentials, exploiting open ports and launching remote execution tools until they reach their objective. Each permissive firewall rule, flat subnet or legacy-sharing channel widens traversal options, enabling rapid spread, deeper persistence and markedly higher recovery costs. An environment designed to minimise attack surface, particularly through microsegmentation, sharply constrains lateral movement options and forces adversaries to expend more resources for each hop.
Command and Control (TA0011)	Even after gaining a foothold, adversaries must establish outbound channels to coordinate operations. Adversaries will seek to abuse unused ports, legitimate outbound protocols such as TCP and DNS, or rogue cloud services so that their malicious tools can receive instructions and exfiltrate data without triggering any security controls. Shrinking the outward-facing surface by blocking unused egress ports, disabling legacy remote-access protocols and applying strict proxy allow-lists limits the transmission paths available for command-and-control traffic. Organisations that have reduced their attack surface will have fewer options for adversaries to establish connections back to their control infrastructure, meaning there is a higher likelihood that beaconing attempts are forced into monitored channels, where they can be detected and disrupted early.

Foundation 7: Resilient networks

MITRE ATT&CK Tactic	Alignment to Foundation 7: Resilient networks
Discover (TA0007)	Attackers enumerate network topology to identify high-value hosts, services, shares and trust relationships. Resilient networks employ strong segmentation, software-defined per-hop encryption and identity-aware routing, which deliberately obscure internal topology and require re-authentication at each zone. Organisations with resilient networks slow attacker mapping efforts and increase the noise generated. Unsolicited enumeration packets will hit dead ends or trigger alerts, improving defender detection.
Lateral Movement (TA0008)	Adversaries escalate by abusing shared authentication tokens, exploiting native remote-management tools and manipulating permissive routing policies to leap across network boundaries toward high-value assets. Poor segmentation and weak route controls allow such movements to proceed with minimal friction, letting attackers expand their operational scope and compromise multiple business functions before detection. Resilient network designs impose strong segmentation, cryptographic route verification and identity-aware tunnels, ensuring that each east-west connection must re-authenticate, thereby isolating breaches, slowing intruders and giving defenders precise choke points for containment.
Command and Control (TA0011)	To manage compromised hosts, adversaries may establish covert outbound channels over web protocols or non-standard ports to communicate with external servers while blending traffic to appear legitimate. Unrestricted egress rules and limited anomaly detection allow these channels to blend with legitimate traffic, providing uninterrupted host control and enabling silent data theft. Resilient networks enforce least-privilege outbound access, apply break-and-inspect gateways and deploy behaviour-based analytics at every exit point. By enforcing strong segmentation, only verified identities can communicate across boundaries, curtailing covert tunnels and letting defenders terminate malicious flows without disrupting vital services.
Impact (TA0040)	Attackers may try to disrupt services, encrypt data, hijack or destroy infrastructure for strategic gain. Destructive actions can cascade across dependent systems, halt revenue streams, threaten safety and complicate already fragile recovery efforts. This may result in organisations facing regulatory fines, contractual penalties and long-term brand erosion. Networks engineered for resilience through redundant paths, automated fail-over and immutable infrastructure components will absorb or confine destructive actions, preserve critical functions and enable rapid recovery, thereby diminishing the adversary's intended impact.

Foundation 8: Secure-by-Design

MITRE ATT&CK Tactic	Alignment to Foundation 8: Secure-by-Design
Resource Development (TA0042)	Threat actors craft malicious tooling that exploits common coding flaws and unsafe defaults. Threat actors may produce malicious software components disguised as valid open-source or proprietary software components. Software built with secure-by-design practices that incorporate actions such as threat modelling, secure coding practices, SBOMs and secure build pipelines with automated security testing reduce exploitable weaknesses, forcing adversaries to invest more effort in bespoke or zero-day capabilities.
Execution (TA0002)	During exploitation phases, adversaries trigger code execution by exploiting memory-safety bugs, injecting commands through unvalidated input, abusing deserialisation routines or loading signed-but-malicious plugins, thereby running attacker code inside trusted processes. Successful execution grants local control, siphons credentials, changes logic and provides a springboard for lateral movement, amplifying both scale and speed of compromise. Secure-by-design enforces least-privilege, leverages memory-safe languages, code signing, integrates rigorous input validation, and enables application allow-listing, dramatically shrinking exploitable surfaces and ensuring exploited components inflict minimal damage.
Privilege Escalation (TA0004)	Adversaries will seek opportunities to elevate their privileges by hunting for exploitable coding oversights, such as improper access controls, which can enable an attacker to gain root or system-level control. Elevated privileges let attackers disable security tools, access protected data and pivot deeper, greatly increasing an adversary's potential impact, as well as their ability to persist and complicate containment. Secure-by-design principles require strict privilege separation, zero trust authentication and authorisation, mandatory access controls and detailed security reviews, sharply reducing escalation avenues and keeping blast radius small.
Defense Evasion (TA0005)	Once operating, adversaries may seek to hide their presence by manipulating processes, sideloading look-alike dynamic link libraries (DLLs) or tampering with logs. These tactics obscure telemetry, frustrate forensics and lengthen dwell time, increasing financial impact and recovery complexity. Secure-by-design embeds tamper-evident logging, validates run time integrity with evidence and implements self-healing protections, making evasion attempts noisy, short-lived and far easier to detect and reverse.

Foundation 9: Comprehensive validation and assurance

MITRE ATT&CK Tactic	Alignment to Foundation 9: Comprehensive validation and assurance
Collection (TA0009)	Adversaries gather sensitive data from all types of data storage to prepare for misuse. Exposure of large-scale aggregation raises legal risks, erodes customer trust and complicates incident response, requiring defenders to trace exactly what left the environment and alert those who are affected. Robust assurance frameworks enforce data classification, least-privilege access and continuously audit settings and data handling controls, limiting the amount of accessible data and flagging unusual bulk-collection activities for rapid investigation.
Exfiltration (TA0010)	After collection, attackers attempt to compress, encrypt and funnel data outside the environment often through misconfigured firewall rules, overlooked cloud storage solutions, covert channels or bulk transfers. Successful exfiltration causes costly breaches, regulatory penalties and public exposure, while encrypted tunnels, high-volume transfers and obfuscated file names negatively impact detection systems and hinder accurate breach scoping for responders. Configuration assurance programmes that enforce data-loss prevention controls, strict outbound encryption policies and regular baseline auditing increase the organisation's ability to detect and often block exfiltration attempts outright.
Impact (TA0040)	Threat actors may seek to corrupt, ransom or destroy data to achieve strategic objectives. Destructive activities extend downtime, inflate recovery costs and may incur legal liability when vital public or private services are disrupted. Configuration validation and assurance –that continually assesses backup schedules, immutability control, privilege boundaries and fail-over settings against approved baselines – are best prepared to respond either through automatic remediation or raising real-time alerts the moment a destructive activity is expected to occur. This allows for swift remediation and reduction of impact.

Foundation 10: Continuous and actionable monitoring

MITRE ATT&CK Tactic	Alignment to Foundation 10: Continuous and actionable monitoring
Execution (TA0002)	Malicious processes often create anomalies in command-line usage, parent–child process relationships or binary provenance. This covert head start lets attackers quietly establish footholds, siphon credentials, manipulate data and embed persistence, expanding the compromise's footprint before responders begin to suspect malicious activity. Continuous, real-time monitoring using high-fidelity signals with automated alerting identifies these anomalies quickly, enabling defenders to isolate affected systems before attackers can progress to further tactics.
Persistence (TA0003)	Attackers can add startup and scheduled tasks, illegitimate accounts and malicious services to survive reboots and credential changes. Persistent implants provide a long-term vantage point for further exploitation and complicate eradication because defenders must locate every mechanism to fully evict the intruder. Constant monitoring of configuration baselines, privilege changes and service-creation events highlights unauthorised persistence mechanisms, allowing security teams to remove them promptly and restore trusted states.
Defense Evasion (TA0005)	Adversaries can disable and clear logs, uninstall agents, alter security configurations or hijack trusted processes to remain hidden. Effective evasion prolongs adversary dwell-time, lets attackers steal more data, and erases evidence investigators need, which in turn raises recovery costs, regulatory exposure and reputational damage. Continuous monitoring measures control-plane health, validates agent integrity, watches for abrupt log cessation and deletion, and detects anomalous behaviour, ensuring that sabotage of security instrumentation itself becomes the signal that triggers rapid containment.
Exfiltration (TA0010)	After collecting data, adversaries attempt to smuggle it out via bulk transfers, staged archives or covert channels. Continuous monitoring that inspects flow volume, data-loss-prevention triggers and protocol anomalies can flag exfiltration in real time, quarantine the responsible sessions and give incident responders clear forensic evidence. Streams of high-volume egress logs, proxy events and storage telemetry all provide critical data that can produce alerts through the analyses platform. Immediate visibility allows for throttling and termination, turning a potential breach into a contained event with minimal data loss.

Supplementary information

The *Information Security Manual* is a cyber security framework that organisations can apply to protect their systems and data from cyber threats. The advice in the Essential Eight Maturity Model prioritises the implementation of controls to mitigate different levels of malicious actors' tradecraft and targeting.

Reference Material:

- UK's National Cyber Security Centre (NCSC-UK) Zero Trust Architecture Design Principles outlines zero trust principles for organisations looking to design and implement a zero trust architecture in an enterprise environment.
- Canadian Centre for Cyber Security (CCCS) A Zero Trust Approach to Security Architecture – ITSM.10.008 provides a description of zero trust security concepts and how organisations can benefit from implementing a zero trust architecture to safeguard their assets.
- NIST SP 800-207, Zero Trust Architecture: NIST's foundational technical publication provides a conceptual framework for zero trust. While not covering all IT, it can be used as a tool to understand and develop a zero trust architecture for an enterprise.
- NIST SP 1800-35, Implementing a Zero Trust Architecture is a series of guides that summarises how the US Government and identified vendors are using commercially available technology to build interoperable, open standards-based zero trust architecture.
- CISA's Zero Trust Maturity Model V2 is designed to provide US federal agencies with a roadmap and the resources to achieve an optimal zero trust environment.
- User; Device; Network and Environment, Data, Application and Workload and Visibility and Analytics is guidance contained in the National Security Agency (NSA) Advancing Zero Trust Maturity Series on zero trust pillars.
- National Security Agency, Embracing a Zero Trust Security Model explains the zero trust security model and its benefits, as well as challenges for implementation.
- US DoD Zero Trust Reference Architecture describes DoD's end-state vision, strategy and framework to strengthen cyber security. It provides technical guidance to evolve existing capabilities with focus on a data-centric security strategy.

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

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 Commonwealth Coat of Arms Information and Guidelines | pmc.gov.au.

For more information, or to report a cyber security incident, contact us:

cyber.gov.au | 1300 CYBER1 (1300 292 371)

ASD AUSTRALIAN SIGNALS DIRECTORATE

ACSC Australian **Cyber Security** Centre