



Australian Government
Australian Signals Directorate

cyber.gov.au

Assessment process guide

Essentials

Contents

Introduction	3
Scope	3
How to use this guide	3
Stages of an assessment	4
Stage 1: Planning and preparation	5
Stage 2: Plan approval	6
Stage 3: Assessment	6
Stage 4: Aggregation and reporting	6
Using the three layers for assessment	7
Why the three layers matter in an assessment	8
Assessing actions	10
The action rating scale	10
Indicators of good practice	11
Assessing domains	13
Scoping a domain	13
Assets that span multiple domains	15
Assessing across the domain	16
When a domain is considered to be addressed	16
Collecting evidence	16
Mapping a control catalogue to the actions	17
Existing evidence as the starting point	17
Strengthening existing evidence	17
The quality of evidence	18
Using tools in an assessment	19
Tools and evidence quality	19
Simulations and active testing tools	19
Using the ATT&CK mapping	20
Informing prioritisation	21
Validating effectiveness	21
Limitations of the ATT&CK mapping	21
Aggregation and reporting of findings	22
Aggregating action ratings to principle ratings	22
Reporting posture	23
The summary view and detailed view	25
Visualising posture and progression	26

Introduction

This guide is a supplemental publication as part of the ASD Essentials series. It explains how organisations can assess the effectiveness of their implementation of an Essentials series chapter.

The purpose of assessment under the Essentials series is to determine whether the security outcomes described by each cyber security principle are being achieved. This is done by assessing the effectiveness of the actions that support each principle.

This guide supports organisations to:

- understand their progress by assessing the applicable principles and actions in each domain and tracking that position over time
- identify and plan remediation by using assessment findings to identify gaps and develop the actions and milestones needed to address them.

Scope

This guide describes an assessment process that can be applied to an Essentials series chapter. The principles, actions, domains and assessment requirements are defined in the applicable chapter and supporting guidance.

This guide supports organisations assessing their own cyber security environment. It does not replace any regulatory or compliance requirements that apply to an organisation. Organisations should consider relevant regulatory requirements when planning and conducting an assessment.

The guide is not intended for assessing the cyber security posture of suppliers or third parties in the cyber supply chain. Guidance for organisations who intend to use the ASD Essentials series to assess their suppliers and third parties will be provided by ASD in the future.

How to use this guide

This guide follows the sequence used when assessing an organisation's implementation of an ASD Essentials series chapter. It begins with the stages of an assessment and the concepts that support assessor judgement, including the three framework layers and IGPs. It then covers scoping and assessing domains, collecting evidence, using assessment tools and the MITRE ATT&CK mapping, and aggregating and reporting assessment results.

Use this guide alongside the relevant Essentials series chapter, the *Introduction to the ASD Essentials series* and the *Implementation planning guide*. These publications provide the framework context and implementation guidance that support the assessment process.

Stages of an assessment



While the approach to an assessment will depend on the scope, size and complexity of the environment, the same stages always apply.

Stage 1 Planning and preparation

The assessor agrees on the scope, engages relevant stakeholders and develops an assessment plan.

Stage 2 Plan approval

The assessment plan is approved before the assessment begins.

Stage 3 Assessment

The assessor conducts an assessment for each domain in scope, rating the applicable actions against assets in that domain and considering the intent of the supporting principle.

Stage 4 Aggregation and reporting

The assessor aggregates action ratings to principle ratings and reports the findings.

The following sections provide detailed guidance for each stage.

Stage 1: Planning and preparation

Thorough planning supports a reliable assessment. Before the assessment begins, the assessor should establish what will be assessed, how it will be assessed and what evidence will be support the assessment.

Agreeing on the scope

The assessor should agree which domains and assets are in scope for the assessment, taking into account the organisation's progress in implementing the relevant Essentials series chapter and it's assessment and reporting needs.

The assessment may include, one, several or all domains. Domains do not need to be assessed in priority phase order. Any domains or assets excluded from the assessment should be documented, including the reason for their exclusion.

Engaging stakeholders

The assessor should identify stakeholders with a role in the assessment and engage them as appropriate. Depending on the organisation, these may include system owners, the security team, infrastructure and platform owners and any managed service providers or third parties who operate parts of the environment in scope. Early engagement establishes the points of contact, confirms who can authorise assessment activities and identifies any operational, contractual or security constraints that may affect how the assessment is conducted.

Where parts of the environment are operated by a third party, the assessor should establish early whether the necessary access and evidence can be obtained, as this often takes time to arrange and can otherwise delay the assessment.

Developing the assessment plan

Having agreed on the scope, the assessor should develop an assessment plan that sets out how the assessment will be conducted. The plan should address:

- the domains and assets in scope, and the principles and actions applicable to each
- the assessment methods to be used for each action, and any scripts, tools or technical testing required
- the existing evidence the organisation holds, and how it will be used as the starting point for the assessment
- access required to systems, data sources, accounts, documentation, personnel and facilities
- sampling methodology, where an action will be assessed across a representative selection of assets rather than every scoped asset within the domain
- how evidence will be collected, handled, stored and protected, including retention or disposal requirements after the assessment
- where the assessment report will be developed and how it will be stored and tracked
- the timeframe, milestones and key dependencies.

The assessment plan should be detailed enough for the organisation to understand what the assessment will involve and identify any activities that require authorisation before they proceed.

Stage 2: Plan approval

The assessment plan should be reviewed and approved by an appropriate authority within the organisation before the assessment begins. Approval confirms that the organisation accepts the scope and the assessment methods, including any proposed technical testing or use of tools and scripts, and that the necessary access and resources will be available.

Approval ensures that the activities carrying operational risk are authorised in advance and conducted within agreed limits. It also confirms that the stakeholders whose systems are in scope are aware of the assessment, and establishes a shared, documented basis for conducting the assessment and interpreting its findings.

If the assessment plan changes materially during the assessment, the changes should be agreed with the organisation and the approval updated accordingly. This includes changes to scope or where additional testing becomes necessary. Any limitation or constraints that affect the assessment, such as access that could not be obtained or testing that could not be performed, should be recorded for inclusion in the assessment report.

Stage 3: Assessment

With the assessment plan approved, the assessor conducts the assessment for each domain in scope. Domains are assessed individually and may be assessed in any order or in parallel, consistent with the agreed scope.

For each domain, the assessor confirms the assets and functions in scope, drawing on the organisation's existing analysis where appropriate. The assessor then works through each principle applicable to the domain and rates each applicable action against the assets in scope. IGPs can be used to support consistent assessor judgement. Each action is judged on coverage, effectiveness, consistency, currency and evidence quality.

The assessor records a rating for each applicable action in each domain, together with the evidence relied on and, for any action rated below **fully effective**, the specific gaps identified. Where the assessment plan allows sampling, the assessor conducts the assessment on a representative selection of the domain's assets rather than every asset in scope. The assessor documents the sample used and the basis for its selection.

Stage 4: Aggregation and reporting

Once the actions have been assessed, the assessor aggregates the action ratings and reports the findings.

Action ratings are aggregated to produce a rating for each applicable cyber security principle in each domain. Applying the weakest link rule, results are reported by priority phase and domain. The assessment report must include:

- a summary view, which shows progression across the domains
- a detailed view, which identifies the specific cyber security principles and actions driving each rating, and the remediation in progress.

Any limitations that affected the assessment are recorded in the report. Detailed guidance for this stage is set out in the [Aggregation and reporting of findings](#) section of this guide.

Using the three layers for assessment

Assessors do not assess controls in isolation. They consider how the three layers of an Essentials framework interact before determining whether an action has been effectively implemented.

- The **principle** describes the enduring security outcome to be achieved.
- The **action** describes what must be done to achieve that principle.
- The **domain** describes where the action is applied and assessed.

		Domain	Domain	Domain
Principle	Action			
	Action		Rating here	
	Action			

Figure 1. Applying the three framework layers in an assessment

Principles define the cyber security outcomes organisations should achieve. Each Essentials framework identifies the principles relevant to its technology environment and applies them through environment-specific actions and domains.

Actions describe what organisations must do to achieve a principle. They are expressed in technology-agnostic terms so organisations can implement controls appropriate to their environment and risk context.

Domains describe where actions apply. Each Essentials framework defines domains that group assets and functions with similar risk characteristics for that technology environment. They help organisations focus security efforts where they will deliver the greatest reduction in risk, rather than applying the same controls across an entire environment.

Priority phases provide ASD's recommended sequence for addressing domains. They do not change how an individual action is assessed. However, an organisation's implementation progress and agreed assessment scope may mean that some domains have not yet been assessed. These domains are recorded as **not assessed** (refer to [Assessing actions](#) for the four-tier assessment scale).

Assessors should assess each action in the context of the domain to which it applies, rather than assessing the action in isolation. The same action may receive different ratings across domains because the assets and risk context differ. Assessors are responsible for determining whether an action has been effectively implemented for the assets and risks within a domain. IGP's provide structured indicators to support consistent assessment criteria, but they do not replace professional judgement. Assessors should consider IGP's together with the intent of the principle and the context of the domain.

Why the three layers matter in an assessment

Effective implementation depends on all three framework layers working together. The same action may require different implementations across domains because the assets, risk context and likely malicious actor techniques differ.

This differs from a checklist approach: A checklist asks whether a control exists, while an Essentials assessment considers whether an action achieves the intended security outcome for the assets in scope.

When assessing an action against assets in a domain, the assessor should be guided by three questions:

1. **What outcome does the principle require for this domain?**

Consider what the intended outcome means for the assets in scope and the consequences of their compromise.

2. **Does the action achieve that outcome for the assets in this domain?**

Assess how the action has been implemented and whether it delivers the intended outcome for the assets in scope.

3. **Is there sufficient evidence that the action is operating effectively across the domain?**

Consider coverage, effectiveness, consistency, currency and evidence quality. An action that is applied to only some relevant assets, or is enforced inconsistently, does not meet the criteria for a **fully effective** rating. Refer to Assessing the actions for the four-tier assessment scale.




Figure 2. Three questions for assessing an action

Example: Vulnerability management in the identity services domain

The **Vulnerability management** principle aims to reduce the likelihood that known vulnerabilities can be exploited to compromise assets. In the **Identity services** domain, the assets in scope include domain controllers, federation services and certificate authorities. As the compromise of these assets can enable broad, cascading access across the environment, the expected outcome is that vulnerabilities are identified and remediated within timeframes commensurate with their risk.

The assessor then evaluates whether the action to apply security patches or updates within defined risk-based timeframes achieves the intended outcome for the assets in the domain. This involves examining whether security updates are applied to identity infrastructure within the organisation's defined risk-based timeframes, with particular attention to vulnerabilities that could enable privilege escalation, credential theft or unauthorised access to identity systems. The assessment focuses on the effectiveness of the action for the domain's assets rather than on patch management practices across the enterprise.

Finally, the assessor determines whether the action is operating effectively and consistently across the domain. Evidence may include vulnerability reports, patch compliance data, change records and configuration management records. If most domain controllers are patched within required timeframes but federation servers or certificate authorities are routinely delayed, the action has not been applied consistently across the domain and would not meet the criteria for a **fully effective** rating.

Assessing actions

Each action is assessed against each applicable domain using a four-tier, outcome-based rating scale: **not assessed**, **not effective**, **partially effective** and **fully effective**.

The action rating scale assesses whether the intended security outcome of the action is being achieved across the assets in scope, rather than whether a specific control has been implemented or a particular configuration applied.

IGPs support consistent assessment of each action by describing characteristics associated with each rating. IGPs for the relevant ASD Essentials frameworks are provided as an annex to this guide.



This assessment approach is informed by the United Kingdom National Cyber Security Centre's Cyber Assessment Framework, which also uses IGPs to support in the assessment of principles-based guidance.

The action rating scale



Not assessed

Not assessed. The action has not yet been assessed against the domain, or there was insufficient evidence to determine whether the intended security outcome is being achieved. This is a neutral rating that may reflect assessment scope, sequencing or a limitation encountered during the assessment. This rating must remain distinguishable from **not effective** in all reporting.



Not effective

Not effective. The action is absent or implemented in a way that does not achieve the intended security outcome for the assets in scope. An action may be present in some form and still be rated **not effective** if it delivers no meaningful security benefit.



Partially effective

Partially effective. The action is in place and achieves the intended security outcome in part, but gaps in coverage, consistency or enforcement, limit its effectiveness across the domain. Partial activity alone, such as work commencing or a plan existing, is not partial effectiveness.



Fully effective

Fully effective. The action is implemented in a way the demonstrably achieves the intended security outcome across the domain and is supported by sufficient evidence.

Indicators of good practice

IGPs are a structural aid to assessment. For each action, they describe the typical characteristics associated with each rating other than **not assessed**. They support informed judgement and consistency between assessors but are not a checklist or substitute for assessor expertise and judgement.

The IGP format

Each action includes three sets of IGP statements, corresponding to the **not effective**, **partially effective** and **fully effective** ratings. The IGP statements are:

- **Not effective** statements describe when the outcome is not being achieved.
- **Partially effective** statements describe when the outcome is being achieved in part but is limited by gaps.
- **Fully effective** statements describe when the outcome is being achieved consistently across the domain.

The **not assessed** rating is not represented by IGP statements. It is used where the action has not yet been assessed or there is insufficient evidence to determine whether the intended security outcome is being achieved.

Action: Record all administrative or management actions that modify systems, configurations or controls with the details of the user, timestamp, and target.		
 Not effective • Administrative actions are not logged or only partially logged. • Logs do not capture sufficient detail to identify who made a change, when and what to target.	 Partially effective • Logging is enabled but does not cover all administrative platforms. • Logs capture the administrator identity and action but lack detail on the target or change nature.	 Fully effective • Comprehensive audit logging to capture all administrative or management actions is enabled on all platforms. • Logs capture administrator identity timestamp, target object, and nature of change.

Figure 3. Example IGP statements

Rating an action based on IGPs

The assessor identifies the IGP statements that best describe the organisation's implementation for the assets in scope. Unless expert judgement supports a different rating, as described in the following section, the lowest applicable IGP level determines the action rating:

- If any **not effective** statement applies, the action is rated not effective
- If all applicable **fully effective** statements apply, and no **partially effective** or **not effective** statements apply, the action is rated **fully effective**
- In all other cases, the action is rating **partially effective**

This rule mirrors the way principle ratings are determined from action ratings (see [Aggregation and reporting of findings](#)). The IGP statements for an action describe different aspects of the same security outcome. A shortfall in one aspect may limit the overall outcome, so the rating generally reflects the weakest applicable aspect.

Supporting expert judgement

The IGPs are intended to inform expert judgement, not replace it. They describe the indicators an assessor would expect to observe at each rating, but they cannot anticipate every environment, technology or compensating arrangement. The assessor remains responsible for determining whether the action achieves its intended security outcome for the assets in scope, taking into account the intent of the principle and the context of the domain.

Several points follow from this:

- **Assess the outcome achieved.** An implementation does not need to match an IGP statement exactly where evidence demonstrates that the action achieves its intended security outcome. An alternative but effective implementation is still **fully effective**.
- **Partial activity is not partial effectiveness.** A partially effective rating must reflect a genuine security benefit, not simply that work has commenced or a control exists in some form. An action that has been deployed but delivers no security benefit to the assets in scope is **not effective**, regardless of the effort involved.
- **A single not effective statement is normally decisive.** A material failure represented by a **not effective** statement will normally result in a not effective rating. Where a compensating arrangement demonstrably achieves the action's intended outcome, the assessor may assign a higher rating if the rationale and supporting evidence are documented.
- **Interpret IGPs in the context of the domain.** The same IGP statement may set a different practical threshold across domains because the assets, risks and consequences of compromise differ. The IGPs should always be read together with the three layer logic described in [Using the three layers for assessment](#).

Where an assessor departs from the IGP statements based on expert judgement, the rationale should be documented so the rating can be understood and reviewed. This should include the outcome achieved, evidence relied on and any compensating arrangements.

Assessing domains

Each domain in the relevant Essentials framework is assessed independently. Domains can be assessed in any order, in parallel or selectively, based on the agreed assessment scope. This differs from ASD's recommended implementation sequence, which uses priority phases to indicate which domains should be addressed. Each domain must be scoped and assessed against its own assets and risk context so that the results accurately reflects the posture of that domain.

Assessors must identify the assets and functions in a domain before assessing the applicable actions. A complete and accurate scope helps ensure the assessment results reflect the organisation's cyber security posture. An incomplete scope can produce misleading results.

This section explains how to scope a domain and assess applicable principles and actions. Detailed guidance on aggregating and reporting assessment results is provided in the [Aggregation and reporting of findings](#) section of this guide.

Scoping a domain

Domains describe where actions apply within an Essentials framework. Scoping identifies the assets and functions relevant to the domain so the applicable actions can be assessed against them.

Where possible, assessors should use the organisation's existing domain scope rather than starting from scratch. Many organisations will have already completed this work when implementing the relevant Essentials framework. More information on scoping domains during implementation is available in the *Implementation planning guide*.

Using the organisation's existing scope reduces effort, keeps the assessment aligned with how the domain is managed and allows the assessor to focus on validating the scope rather than rebuilding it. The assessor should confirm that the existing domain scope is accurate, complete and current using sources such as asset inventories, system registers and architecture documentation. Minor gaps should be reflected in the assessment scope. Significant gaps that call the completeness of the scope into question should be recorded as findings and may require broader asset discovery or asset management activity before the assessment results can be relied on.

When scoping a domain, assessors should consider the asset categories defined by the relevant Essentials framework. For the *Essentials for enterprise IT* framework, these include infrastructure, operating systems, applications, identities, credentials and data.

Table 1 shows examples of the asset categories described in the *Essentials for enterprise IT* framework.

Table 1: Asset categories and examples

Asset category	Examples
 Infrastructure	servers, workstations, laptops, mobile devices, network devices, storage systems, virtual machines, hypervisors, container hosts and cloud-hosted compute, storage and networking resources
 Operating systems	desktop, server, mobile, network-device, appliance and embedded operating systems, including operating system images and versions
 Applications	commercial and internally developed applications, software-as-a-service applications, middleware, database platforms, application programming interfaces, container images, software libraries and development platforms
 Identities	user identities, administrative identities, service accounts, machine identities, workload identities, application identities and other non-person identities
 Credentials	passwords, authentication keys, secrets, tokens, certificates and other authentication material associated with an identity or service
 Data	business records, data sets, databases, file repositories, object storage, source code, configuration repositories, backups and archives

Not every asset category will be relevant to every domain. The asset categories provide a prompt for scoping, rather than a requirement that each domain contain assets of every category.

Assets that span multiple domains

An asset does not always belong to just one domain. Some assets support multiple domains because they perform different functions or contain components relevant to more than one area.

When an asset supports more than one domain it should be included in each relevant domain. Each domain is assessed separately, so the effectiveness of actions applied to the asset may be rated differently depending on the domain and its risk context.

Assessment evidence for a shared asset can be reused across domains where relevant, but the assessor should evaluate that evidence against the applicable actions and risk context of each domain. Figure 4 illustrates how an asset can support more than one domain.




Figure 4. Examples of an asset spanning multiple domains

Assessing across the domain

Once the domain is scoped, the assessor works through each principle applicable to the domain and rates the applicable actions against the assets in scope. IGPs can be used to support the assessor's judgement.

Not every principle or action applies to every domain. The relevant Essentials framework identifies which actions to apply to each domain. Only actions mapped to the domain are assessed for that domain.

For each applicable action, the assessor examines four assessment dimensions to ensure effectiveness:

- **Coverage** – whether the action is applied to all relevant assets within the domain.
- **Consistency** – whether the action is applied consistently across those assets.
- **Currency** – whether the evidence is recent enough to reflect the current implementation and operating intent of the action.
- **Evidence Quality** – whether the evidence sufficiently reliable and relevant to demonstrate that the intended outcome is being achieved.

Gaps in coverage or consistency may prevent an action from being rated **fully effective**. The assessor should consider the extent and impact of the gaps when determining the appropriate action rating. Where an action is rated **partially effective**, the assessor should document the specific gaps so that remediation actions and timeframes can be developed.

When a domain is considered to be addressed

A domain is considered to be addressed when all applicable actions are operating as intended and are rated as **fully effective**. A domain that has not yet been assessed is recorded as **not assessed** across its applicable principles, reflecting sequencing rather than failure.

Because each domain is assessed individually, its result stands on its own. Results in one domain do not carry across to another. Each domain must be assessed against its own assets and risk context.

Collecting evidence

ASD Essentials frameworks describe security outcomes and actions in technology-agnostic terms rather than prescribing specific controls. Organisations can achieve these outcomes using controls, processes and technologies appropriate to their environment. Existing control catalogues or frameworks can provide sources of assessment evidence, but the assessment focuses on whether the intended security outcome is being achieved rather than on the controls themselves.

This has direct implications for assessment. Many organisations will already hold evidence that can be used to assess the effectiveness of actions. Assessors should identify and use relevant existing evidence before generating new evidence.

Mapping a control catalogue to the actions

Organisations can map existing controls to Essentials actions to demonstrate how they achieve the required outcomes. The ISM is one source of controls, but other frameworks and control catalogues can also be used. Control mapping should be used as a starting point for implementation and assessment activities rather than as evidence of that the outcome has been achieved. While mapping can help identify relevant controls and evidence sources, organisations must still demonstrate through assessment and supporting evidence that the action's intended security outcome is being achieved and remains effective.

The assessor determines whether the controls implemented by the organisation achieve the outcome required by the action. The assessment focuses on the outcome rather than alignment with a specific control catalogue, so different organisations can use different controls and still achieve the same action rating.

Existing evidence as the starting point

An organisation that has implemented and assessed controls under another control catalogue will likely hold evidence, such as configuration baselines, assessment reports, control test results, audit findings and risk registers. The assessor should identify which evidence is relevant and current to each applicable action and reuse it where appropriate.

If existing evidence remains representative of the current environment, reflects the controls and processes in operation at the time of assessment, and is reliable, consistent and sufficiently comprehensive across the domain, no further evidence is required. Where evidence remains relevant but is incomplete or applies to only part of the domain, it may reduce the additional assessment activity required.

While existing evidence provides a valuable foundation for assessment, it should not be relied upon without qualification. Assessors should be aware of three considerations.

- **Deployment is not effectiveness.** Evidence that a control is configured or enabled does not prove it is achieving the intended outcome. The assessor should look for evidence of **fully effective** operation across the domain.
- **Evidence must be current.** Older evidence may no longer reflect the current environment. Assessors should determine whether existing evidence remains relevant and seek more recent evidence where necessary.
- **Controls may only partially satisfy an action.** A mapped control may address only part of an action's outcome. Identifying and assessing any remaining gaps is a key part of the assessment.

Strengthening existing evidence

Where existing evidence demonstrates deployment but not sustained effectiveness, it can be supplemented with operational evidence from technical validation activities, such as red team exercises, purple team activities, control validation testing, breach and attack simulation and penetration testing. The [Using the ATT&CK mapping](#) section explains how techniques mapped to an action can inform technical validation activities.

The quality of evidence

Not all evidence supports a rating equally. When assessing an action, the assessor should consider the quality of the evidence available and seek the highest quality that is reasonably practicable. Figure 5 illustrates four broad levels of evidence quality.




Figure 5. Evidence pyramid

Four broad levels of evidence quality apply:

- **Excellent evidence** provides direct evidence that the action works as intended in the live environment.
- **Good evidence** shows the action operating across the assets in scope, including evidence of coverage, consistency and ongoing effectiveness.
- **Fair evidence** indicates the action is likely operating as intended but does not directly demonstrate this across the domain.
- **Poor evidence** shows or asserts that the action exists but does not directly demonstrate that it is operating effectively or achieving its intended outcome.

These evidence levels are a guide, not a rule. A single strong piece of evidence may not be enough to demonstrate an outcome across an entire domain, while multiple pieces of good evidence may provide a stronger overall picture. Assessors should consider all available evidence and the risks associated with the domain.

Lower quality evidence may still contribute to an assessment, but its limitations should be considered when determining the action rating. Where the available evidence is insufficient to determine whether the intended outcome is being achieved, the action should be rated **not assessed** and the evidence limitation documented. Where the evidence demonstrates that the intended outcome is not being achieved, the action should be rated **not effective**.

Using tools in an assessment

Tools can help assessors collect information, identify gaps and confirm whether an action is working as intended. This can improve the quality and coverage of assessment evidence.

This guide does not recommend specific tools. Different organisations use different technologies, so the most suitable tool will depend on the environment and the action being assessed. Assessors should choose appropriate approved tools in accordance with the assessment plan and consider native capabilities before introducing additional products. The tool, version, configuration permission, licensing constraints and known limitations should be recorded. Tool-generated evidence should be critically assessed for incomplete or stale information, misconfiguration, access restricting and false-positive or false-negative results.

Tools and evidence quality

A tool's evidence should be assessed the same way as any other evidence. Tools can provide strong evidence when they directly test whether an action is achieving its intended outcome, or they can provide supporting evidence by showing configuration and operational status. The quality of the evidence depends on what the tool demonstrates, not on the fact that a tool was used.

Like any evidence, results from a sample of assets do not automatically demonstrate an outcome across the entire domain. Assessors should consider tool output alongside other evidence, taking into account coverage, consistency, currency and evidence quality, as well as the risk context of the domain.

Simulations and active testing tools

Active testing can provide strong evidence by directly testing whether an action works as intended in the live environment, rather than relying on configuration data alone. Examples include attempting to run unapproved software or simulating techniques that the action is intended to prevent or detect.

Because active testing generates activity in live systems, it can introduce operational risk. Testing should be documented in the assessment plan and approved by the appropriate authority before it begins, including the target systems, scope and agreed limits.

Active testing should be conducted by personnel with the skills or experience required to manage the associated risks. Where appropriate capability is not available within the organisation, suitably qualified external capability may be required.

Using the ATT&CK mapping

Most actions in the framework are mapped to one or more MITRE ATT&CK techniques. These techniques describe behaviours used by real-world malicious actors. Actions mapped only where they directly help prevent, detect or limit a technique.

For assessment purposes, the ATT&CK mapping is a guide rather than a scoring mechanism. Actions are still rated against their intended security outcomes. The mapping can help assessors prioritise assessment and remediation activities, identify opportunities to validate implementation effectiveness, and highlight potential gaps in detection coverage. Figure 6 shows how ATT&CK mapping can inform assessment prioritisation and validation.




Figure 6. Using ATT&CK to support assessment

Informing prioritisation

The ATT&CK mapping can help assessors and organisations prioritise assessment and remediation activities based on the threats relevant to the environment. Comparing observed malicious actor techniques with mapped actions can help identify gaps that warrant earlier attention. This can be particularly useful when planning remediation for actions rated **partially effective** or **not effective**. The mapping may inform the urgency of remediation, but it does not change the action rating.

Validating effectiveness

The mapping can also support technical validation. Because actions are mapped to techniques that they are intended to disrupt, assessors can design testing that exercises relevant behaviours and observe whether the implemented controls prevent or detects them in the live environment. This can provide high-quality evidence because it tests whether the intended outcome holds in practice rather than inferring effectiveness from configuration alone.

Example: Trustworthy software

An organisation that has rated an action under the Trustworthy software principle as **fully effective** can use ATT&CK informed testing to validate that rating. For example, testing could attempt to execute unsigned binaries or use unapproved interpreters to confirm whether the controls prevent or detect the activity as intended. The results can provide additional evidence about whether the action is operating effectively in the domain.

Testing of this kind can introduce operational risk and must be planned and approved by an appropriate authority before it is conducted. The techniques or behaviours to be tested, target systems and testing limits should be documented in the assessment plan.

Limitations of the ATT&CK mapping

ATT&CK mapping informs assessment and decision-making but does not replace assessor judgement. Assessors should be aware of the following limitations.

- **ATT&CK evolves continuously.** The MITRE ATT&CK knowledge base changes over time. New techniques are added, existing techniques are revised and relationships with mitigations may change. ASD's mapping will be updated periodically, but organisations making high-impact decisions should also consider the current MITRE ATT&CK knowledge base.
- **Mapping does not prove implementation effectiveness.** The mapping is necessarily generalised. An action mapped to a technique does not guarantee that the organisation's implementation will mitigate that technique in practice. Effectiveness depends on how the action is implemented in the specific environment and should be assessed using the outcome-based approach described in this guide.
- **Threat intelligence is incomplete.** Cyber threat intelligence is inherently incomplete. Malicious actors adapt their techniques, and relying only on previously observed activity may understate emerging or unobserved threats. ATT&CK mapping should therefore inform prioritisation alongside the broader risk management process.

Aggregation and reporting of findings

This section explains how action ratings are aggregated into principle ratings and how results are reported. It also describes reporting formats used to present an organisation's cyber security posture to executives, governing bodies and stakeholders.

Aggregating action ratings to principle ratings

Action ratings are aggregated to produce a single rating for each principle within each domain. A principle is rated separately in each domain to which it applies because its effectiveness depends on the assets and risk context of that domain.

Within a domain, a principle may have several applicable actions. The principle rating is determined by the lowest applicable action rating.

- **Not effective** – one or more applicable actions are rated **not effective**
- **Partially effective** – all applicable actions have been assessed, none are rated not effective and at least one is rated **partially effective**
- **Fully effective** – all applicable actions are rated **fully effective**

The weakest link rule reflects that the actions under a principle work together to achieve its intended security outcome. A material gap in one action can therefore limit the effectiveness of the principle within that domain. The principle rating reflects the lowest assessed action rating.




Figure 7. Aggregating action ratings to a principle rating

Example:

The **Identity, credentials and access management principle** has six applicable actions under the Identity services domain. Four actions are rated **fully effective** and two are rated **partially effective**. The principle is therefore rated **partially effective** in that domain. The same principle is rated separately in the Boundary services domain. If all applicable actions in that domain are rated **fully effective**, the principle is rated **fully effective** there.

Where the applicable actions for a principle have not been sufficiently assessed to determine a principle rating, the principle is recorded as **not assessed** for that domain. Where no applicable actions have been assessed within a domain, the applicable principles are recorded as **not assessed** and the domain is reported accordingly.

Reporting posture

Results are reported by priority phase and domain, showing the number of principles at each rating. This format allows progression across multiple principles to remain visible even where other principles are still **partially effective** or **not effective**.

A percentage is calculated for each domain to summarise progress. Each principle rating is assigned a score: **not effective** is scored 0, **partially effective** 0.5 and **fully effective** 1. The sum of the principle scores is divided by the number of applicable principles in the domain.

Example:

In the Visibility and response domain, an organisation assesses **5 principles** as **fully effective** ($5 \times 1 = 5$), **4** as **partially effective** ($4 \times 0.5 = 2$) and **1** as **not effective** ($1 \times 0 = 0$). The score is 7, divided by 10 principles, giving **70%**.

Table 2 illustrates summary reporting for an organisation working towards priority phase 2.

Table 2: Example summary assessment results

Priority Phase	Domain	Fully effective	Partially effective	Not effective	% effective	Not assessed
All	Governance and culture	7	0	0	100%	0
All	Visibility and response	5	4	1	70%	0
1	Boundary services	9	7	0	78%	0
1	Identity services	9	5	2	72%	0
2	Critical processes and services	12	4	0	88%	0
2	Privileged environments	7	6	3	63%	0
3	Business systems	0	0	0	N/A	13
3	User devices	0	0	0	N/A	10

The summary view and detailed view

The assessment report must include both a summary view and a detailed view. The summary view demonstrates progression over time, while the detailed view identifies the principles and actions driving each rating and any remediation activity in progress.

Both views should be available to the intended reporting audience. The summary view alone may obscure outstanding risk, while the detailed view identifies the specific gaps requiring attention. Figure 8 shows how the summary view can be expanded to identify the principles and actions driving a domain result.




Figure 8. Summary and detailed assessment views

Visualising posture and progression

The reporting tables provide the structured detail required for assessment and remediation planning, but they may not be the most suitable format for executive or board briefings, or rapid comparison across domains. Complementary visualisations can present the same underlying data in a format suited to these audiences.

A radar chart, such as the example shown in Figure 9, can provide a high-level view of an organisation's posture across multiple domains or principles. Overlaying results from successive assessments can also show progression over time.

Two views may be useful:

- **a domain-level view** showing results across the domains, suited to executive, board and regulator reporting
- **a principle-level view** showing results for the principles applicable to a single domain, suited to CISOs, security architects and assessors.




Figure 9. Essentials for enterprise IT – domain posture

Visualisations must not be used in isolation. Aggregating ratings into a single point can obscure the underlying detail and individual gaps. Visualisations should always be presented alongside the structured reporting that identifies the principles and actions driving each rating.

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

cyber.gov.au

1300 CYBER1 (1300 292 371)

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, the entity's logo, third party content, 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>

