



Australian Government
Australian Signals Directorate

cyber.gov.au

ASD Essentials overview

Which chapter should I start with?

Essentials

Choosing a chapter

The Essentials series currently comprises chapters for:

- enterprise IT
- cloud
- operational technology.

Organisations may operate environments covered by more than one chapter, and there is no requirement to work through them in a set order. The chapters are designed to be applied independently, in the sequence that best reflects an organisation's environments, obligations and risk. The following questions are intended to help determine which chapter to begin with.




Figure 1. Selecting a chapter

What do you operate?

The starting point is an accurate understanding of what digital environments the organisation operates or has a responsibility to manage. This is not a comprehensive asset discovery exercise, but a sufficient overview of the technology environments in use, the functions they support and the assets they hold.

An organisation may find that only one chapter is relevant, that several chapters apply to distinct parts of its environment, or that a single business function depends on assets falling under more than one chapter. Where environments are interconnected, it is useful to note the points at which they meet, as compromise in one environment frequently provides a pathway into another, and a potential vulnerability that increases the risk.

What are you required to do?

Organisations operate under a range of regulatory, contractual and sector-specific obligations and these frequently direct attention to particular environments. Where an obligation requires an organisation to demonstrate a specific security posture for a defined environment, the chapter covering that environment is a logical starting point. Regulatory drivers should be considered alongside the following risk-based questions rather than in place of them: an obligation establishes a floor, not a priority order.

Where is the biggest gap?

Organisations should consider the cyber security gap between the current state of each environment and the outcomes described by the applicable principles. Some environments that have received sustained investment may require less work than those that have not.

Operational technology environments are a common example. Operational Technology (OT) environments often carry significant technical debt: long asset lifecycles mean systems remain in service for decades, availability and safety requirements constrain patching and change windows, and legacy protocols and unsupported platforms frequently remain in operation because replacing them requires capital investment and operational disruption that has not yet been scheduled. An organisation may find its OT environment has a materially higher risk than its enterprise IT environment, given OT's importance to critical infrastructure assets that supply essential services to the public.

A large gap is not, by itself, a reason to start with an environment. It is one input to consider alongside the consequence and likelihood of compromise.

Where is the greatest risk?

For each environment, organisations should consider two questions about the cyber security risks being managed.

What would the consequence of compromise be? Consequence can differ significantly between environments depending on the criticality to the business function or the impact to the organisation's objectives.

For example, the compromise of an OT environment may affect personnel safety, physical processes or the continuity of an essential service. Compromise of an enterprise IT environment may result in loss of sensitive data, disruption to business operations or loss of the identity infrastructure that underpins access to everything else. Compromise of a cloud environment may expose data held with a provider, or the administrative planes through which services are configured and controlled.

How likely is compromise? Likelihood reflects the probability of the risk occurring, based on the value, vulnerability and exposure of the environment. Environments accessible from the internet, available to a large user base, with opaque access controls or known to be targeted across the organisation's sector, carry higher likelihood than those that are isolated and narrowly accessed.

Environments with both high consequence and high likelihood warrant attention first. Where an environment carries severe consequence but lower likelihood, the position should be documented and reviewed, as the risk changes due to potential variations to connectivity, remote access arrangements and integration decisions.

Enterprise IT or cloud?

Organisations often need to decide whether to start with the enterprise IT chapter or the cloud chapter. This depends on the type of services they use and how much control they have over the underlying technology and infrastructure.

Where an organisation operates on-premises infrastructure, a hybrid environment or infrastructure-as-a-service, the enterprise IT chapter is the ideal starting point. In these models the organisation retains responsibility for the operating systems, applications, identity infrastructure and network architecture that the enterprise IT chapter addresses, and the fact that some of it runs on infrastructure provided by another party which does not materially change what the organisation must do.

Where an organisation operates predominantly platform-as-a-service or software-as-a-service, and requires a clear understanding of which security responsibilities sit with the organisation and which sit with the provider, the cloud chapter is the appropriate starting point. The cloud chapter addresses the shared responsibility model directly, and its actions are written for environments where the organisation configures and governs a service it does not operate.

Most organisations will need both chapters over time. The question is which to begin with, not which to adopt.

Sequencing the work

Each chapter groups its domains into priority phases. The phases set out the order in which ASD recommends the domains be addressed, and reflect where ASD assesses that effort delivers the greatest and rapid reduction in risk.

Priority phases are a sequencing mechanism, not a rating. They do not represent assurance over the effectiveness of controls, compliance against a standard or an organisation's cyber security maturity against a particular class of adversary. An organisation working at an early phase is not operating at a lower standard than one working at a later phase; it is at a different point in the same program of work. The recommended sequence is designed to support organisations to progress through each phase, scaling and hardening their environments as they go, reducing the organisations' overall cyber security risk.

Why this order

Adversary tradecraft

The domains addressed first are those most commonly used by attackers to gain access, and move through an environment towards valuable assets. Strengthening these areas early helps block the techniques attackers rely on most and limits the damage they can cause if they reach other parts of the environment.

Current gaps in control coverage

ASD's intelligence, incident data and partner feedback highlight areas where security controls are often weakest across Australian organisations. Where a gap is both common and frequently exploited, the domains that address it are given higher priority.

Technology and solution opportunities

The sequencing also considers what organisations can realistically implement. Domains with mature, widely available solutions can often be addressed more quickly and with less effort, while also providing a foundation for improvements in other areas.




Figure 2. Progressing through implementation levels

Changing the order

The recommended sequence is a starting position, not a constraint. Organisations are best placed to understand their IT environments, obligations, threat exposure and delivery capacity better than any general guidance can. Organisations should adopt each chapter in the way that best serves them.

There are valid reasons to follow a different sequence than the one recommended. For example, threat intelligence may show that a particular domain is being actively targeted and needs immediate attention, even if it appears later in the sequence. Regulatory or contractual requirements may also require certain capabilities to be implemented first.

Organisations may also have opportunities to address a domain sooner than planned. A technology refresh, system replacement, or existing project may make it more efficient to implement related controls while the work is already underway. In some cases, an organisation may have already made strong progress in the domains addressed early, meaning the greatest reduction in risk can be achieved by focusing on a different area.

If an organisation decides to change the recommended sequence, it should document the reasons for doing so. This helps when reviewing decisions later, explaining the approach to executives, boards or regulators and ensuring the sequence can be adjusted as risks and business needs change.

Regardless of the order chosen, two considerations remain important. First, domains addressed later in the sequence will still need to be implemented. Changing the sequence only changes what is done first, not what needs to be done overall. Second, domains that apply continuously across all levels should be established early and strengthened over time. These domains provide the foundation for the rest of the security program and support the effectiveness of other controls.



Reporting against domains

Regardless of the sequence an organisation follows, domains remain the basis for measuring and reporting its security posture.

Domains are designed around business functions and the assets that support them. This means that gaps can be described in terms of the functions they affect, rather than as isolated control weaknesses. Reporting at the domain level helps executives, boards and regulators understand not only that work remains to be done, but also which parts of the organisation are exposed to risk as a result.

This approach also supports organisations that choose a different implementation sequence. By reporting consistently against domains, an organisation can clearly show what has been completed, what has been deferred, the reasons for those decisions, and the risks that remain. This ensures that flexibility in sequencing does not reduce visibility of the organisation's risk position.

Domains that have not yet been addressed should be reported accordingly. A domain that has not been assessed simply reflects the chosen sequence of work and should be clearly distinguished from a domain that has been assessed and found to have deficiencies.




Figure 3. Example Boundary services domain report

Outsourced services

Organisations often rely on third parties to deliver, host, operate or support the systems covered by the Essentials series. However, outsourcing a service does not transfer accountability for the security outcomes that service is expected to achieve.

When a provider is responsible for delivering all or part of a security requirement, the organisation remains responsible for ensuring it is being delivered and that it meets the required standard. This includes understanding which responsibilities belong to the provider and which remain with the organisation, obtaining sufficient evidence that the provider is meeting its obligations, and ensuring that evidence remains current rather than relying solely on assurances provided when the contract was signed.

If a responsibility is not being met, or there is insufficient evidence that it is being met, the resulting risk remains the organisation's responsibility. The organisation should understand the impact of the gap, determine whether it can be addressed through the provider, compensating controls, or changes to the arrangement, and where it cannot be fully addressed, document and formally accept the residual risk. Accepted risks should be reviewed regularly, as both threats and provider practices can change over time.

The same applies to responsibilities retained by the organisation. If a provider offers a security capability but the organisation has not enabled or configured it, the responsibility has not been fulfilled.

Gaps arising from outsourced services should be recorded and reported against the relevant domains in the same way as gaps in internally managed environments. A domain should not be considered effective, simply because responsibility for it has been assigned to a third-party.



**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>

