Black Shard

TECHNICAL WHITEPAPER

What each Essential Eight maturity level demands

Level 1, Level 2 and Level 3 ask more of the same eight strategies, and the model's own words set the bar at each. The patch clock is the part that moves fastest, and the level a firm holds is the level of its lowest strategy.

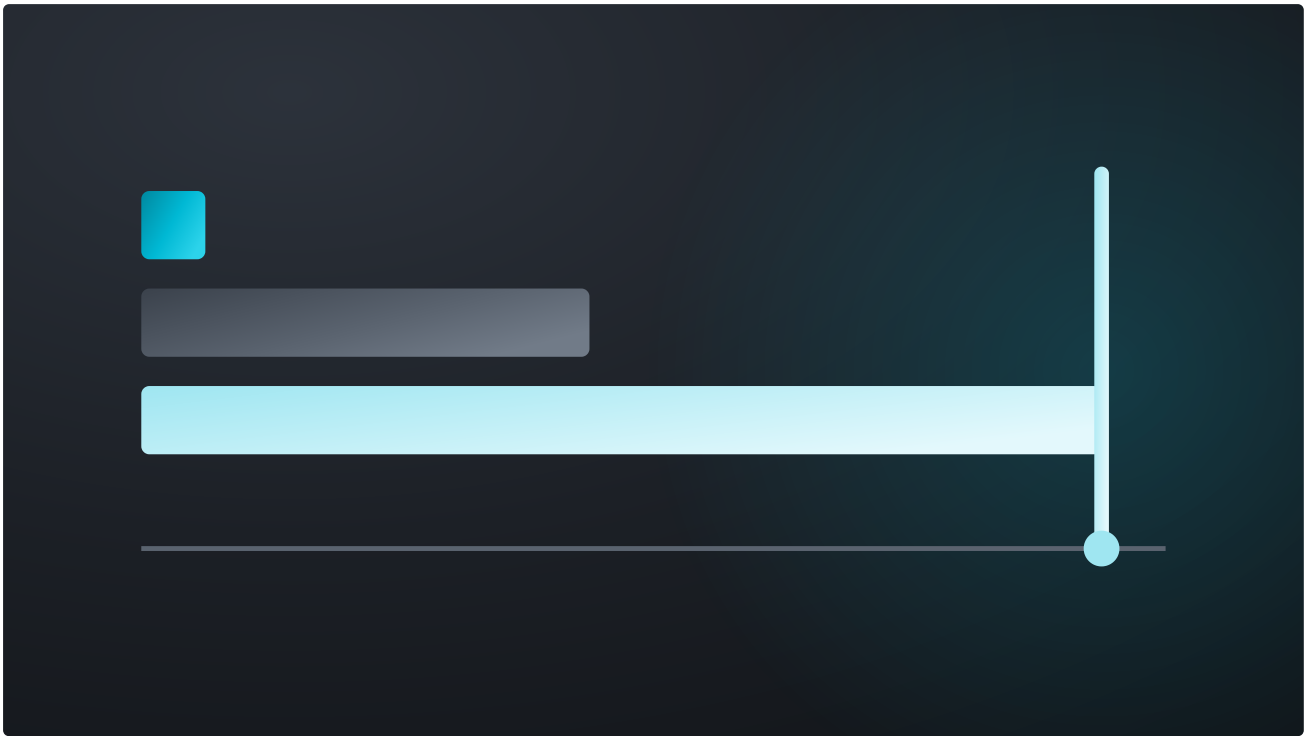

Entity	Black Shard Pty Ltd · ABN 66 696 910 773
Document ID	BS-RPT-07
Classification	Public
Version	1.2
First published	17 March 2026
This edition	16 November 2026
Figures	Model of Nov 2023, read 17 Mar 2026

Executive summary

The Essential Eight Maturity Model of November 2023 defines four maturity levels, Level 0 to Level 3, and rates each of the eight strategies against the tradecraft and targeting the level is meant to defeat. At Level 1, Level 2 and Level 3 alike, a vulnerability in an online service or in the operating system of an internet-facing server or network device is patched within 48 hours of release when the vendor rates it critical or a working exploit exists.

Each level is written against an adversary. Level 1 is set against opportunists using commodity tools on any victim they can reach. Level 2 is set against attackers who invest in a particular target, phish for credentials and work around weak multi-factor authentication. Level 3 is set against adaptive attackers who move on a public exploit the day it appears, steal authentication tokens, hunt privileged credentials and cover their tracks. The requirements at each level are the controls that defeat that adversary and no more.

Online services, internet-facing operating systems and the applications that handle content from the internet (office productivity suites, web browsers and their extensions, email clients, PDF software and security products) are otherwise patched within two weeks of release at every level; the November 2023 edition cut the application window from one month. The clock is the part of the model firms most often misjudge. Six of the strategies are a state a system is in; the two patching strategies carry a stopwatch, and the stopwatch starts when the vendor publishes, whatever day of the week that is.

For a small or mid-sized firm the practical reading is this. Level 1 is a floor every business can reach with the products it already licenses. Level 2 is the level a government buyer or an insurer usually has in mind when the questionnaire says Essential Eight, because it is the first level that demands phishing-resistant sign-in and central logging. Level 3 is for firms that hold information worth a targeted effort, or that supply those who do.

- An organisation is rated at the level of its weakest strategy, so the first pass is every strategy to Level 1 before any strategy goes higher.
- The fastest patch clock applies at every level; what the level changes is how much of the estate it covers.
- Every requirement is read back from the system that enforces it, with its mode and its assignment. A policy document on its own rates Level 0.

Method

The source is the November 2023 edition of the maturity model on cyber.gov.au, published on 27 November 2023 and read on 17 March 2026: the level descriptions, the three appendices that list the requirements for Level 1, Level 2 and Level 3, and the companion page that explains what changed in that edition. Every requirement below is the model's wording, condensed; every window and count printed is the figure the page gives, and nothing is read in from any other edition or guide.

A requirement in the model is a state a system is in. 'Microsoft Office macros in files originating from the internet are blocked' is a policy setting with a scope, and an assessor wants three things for it: the policy as the enforcing console shows it, the mode it runs in, and the users or devices it is assigned to, with a sample of devices showing the setting applied. The method a reader applies in their own estate is the same for all eight. Name the system that enforces the strategy. Export its configuration from that system, never from the design document. List every exception with the date it was granted and the date it closes. Then give the control a check that runs every day and tells someone when the state changes.

The order of attack follows cost and damage. Multi-factor authentication and backups come first: both are policy applied to accounts and storage a firm already has, and between them they stop the stolen password and the ransom. Administrative privileges come next, because every later control is undone by an administrator account that signs in from a daily workstation. The patch clock follows, as a process before a product: who learns of a critical release, who approves the change, and how it is applied inside the window on a weekend. Macro settings and application hardening are a week of policy work. Application control is last, because an allowlist has to admit every legitimate program before it can block the rest, and that takes a period of watching before enforcing.

The levels are cumulative, and an organisation holds the level of its weakest strategy. A firm that reaches Level 3 on seven strategies and Level 1 on the eighth holds Level 1. That rule sets the order of work: every strategy to the target level before any strategy goes beyond it.

The three levels, as the model describes them

The model defines its levels by tradecraft and targeting, and tells an organisation to choose a target level by how desirable it is to an attacker and what a compromise would cost it. Level 0 is the absence of a level: weaknesses in the overall posture that, when exploited, let an attacker of the kind Level 1 describes compromise the firm's data or systems.

Level 1

The model writes Level 1 against 'malicious actors who are content to simply leverage commodity tradecraft that is widely available'. Its examples are a public exploit used against an online service that was never patched, and a sign-in to an online service with a password that was stolen, reused, brute forced or guessed. These attackers look for any victim. They use common social engineering to get a user to weaken a setting or run a program, exploit whatever privilege the account they land in carries, and may destroy data, backups included.

The requirements match. Patch the edge fast and the rest on a calendar, put multi-factor authentication on every service that holds sensitive data, keep administrators on separate accounts and separate environments, block the programs and macros that arrive from the internet, harden the browser, and keep backups that an ordinary user cannot delete.

Level 2

Level 2 is written against 'a modest step-up in capability': attackers willing to invest more time in a target and in the effectiveness of their tools. They phish for credentials and use technical and social engineering techniques to get around weak multi-factor authentication. They are more selective in whom they attack, still conservative in what they spend, and when the account they compromise has no special privileges they go looking for one that does. They may destroy every backup a privileged account can reach.

So Level 2 is where phishing-resistant multi-factor authentication becomes a requirement for users of online services and of systems, where privileged access expires unless it is revalidated, where administration runs through jump servers, where application control reaches internet-facing servers and every folder, and where central logging arrives: multi-factor, privileged access, application control and script events are logged, protected from change, analysed in a timely manner, and incidents are reported to the chief information security officer and to ASD.

Level 3

Level 3 is written against actors who are ‘more adaptive and much less reliant on public tools and techniques’. They exploit older software and thin logging to extend their access and hide, make swift use of exploits as they become public, and get past stronger multi-factor authentication by stealing the token that proves a user has already signed in. Once inside they seek privileged credentials or password hashes, move to other parts of the network and cover their tracks.

Level 3 puts the 48-hour clock on critical vulnerabilities everywhere: in the applications that handle internet content, in the operating systems of workstations and internal servers and network devices, and in drivers and firmware. Level 3 also brings just-in-time administration, secure administrative workstations and the operating system’s credential protection features, extends application control to internal servers and to drivers, allows only signed or sandboxed macros, removes the old script engines, puts the shell into its constrained language mode, and analyses the logs of internal servers and workstations as well as those of the edge.

The model states its own ceiling. Level 3 ‘will not stop malicious actors that are willing and able to invest enough time, money and effort to compromise a target’, and it points an organisation at that point to the rest of the Strategies to Mitigate Cyber Security Incidents and to the Information Security Manual. It also says that no independent assessment is required unless a government directive, a regulator or a contract asks for one, which is exactly where most small firms meet it.

The patch clock

THE PATCH CLOCK, IN DAYS FROM RELEASE

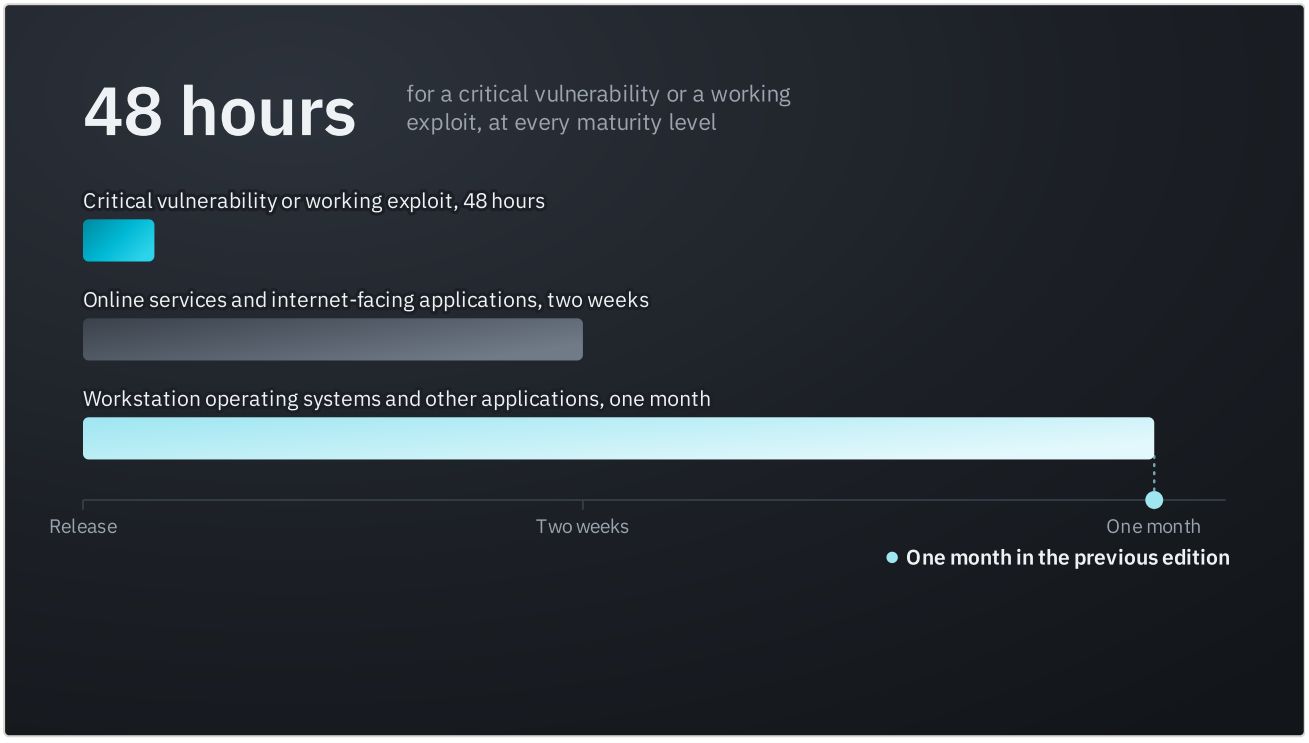

One clock for a critical vulnerability or a working exploit, one for the software that touches the internet, one for the rest; the level decides how much of the estate the fastest clock covers. Source: ASD, Essential Eight maturity model and maturity model changes, November 2023 edition, read 17 March 2026.

At Level 1, Level 2 and Level 3 alike, a vulnerability in an online service or in the operating system of an internet-facing server or network device is patched within 48 hours of release when the vendor rates it critical or a working

exploit exists. Online services, internet-facing operating systems and the applications that handle content from the internet (office productivity suites, web browsers and their extensions, email clients, PDF software and security products) are otherwise patched within two weeks of release at every level; the November 2023 edition cut the application window from one month.

Operating systems on workstations and on servers and network devices that do not face the internet are patched within one month of release, and from Level 2 every application outside the internet-facing set is patched within one month as well. Level 3 puts the 48-hour clock on critical vulnerabilities everywhere: in the applications that handle internet content, in the operating systems of workstations and internal servers and network devices, and in drivers and firmware.

Three clocks, then, and the level decides their reach. At Level 1 the fastest clock covers the internet edge only; from Level 2 the one-month clock gains the rest of the application estate; Level 3 puts the fastest clock on every critical flaw wherever it sits, firmware included. The scanning cadence sits underneath: at least daily against online services and internet-facing systems, at least weekly against the applications that handle internet content, at least fortnightly against everything else, with an automated asset discovery at least fortnightly so the scanner knows what to scan.

The changes page explains why the November 2023 edition looks the way it does. ASD measured how quickly attackers were exploiting new vulnerabilities and shortened the clocks where that speed does the most damage: the critical window at every level, and the internet-facing application window from one month to two weeks. To pay for that, the window for operating systems on workstations and internal devices was loosened from two weeks to one month. The edition also added drivers and firmware to Level 3. A firm whose patch schedule still follows the previous edition meets the old windows and misses the new ones on the internet edge.

In practice the fastest clock is a process question before it is a tooling question. Someone has to learn within hours that a critical release exists for the edge, someone has to approve the change, and the change has to be applied inside two days, including the two days that fall on a weekend. Firms that hold the window have written down who does each step and have a standing approval for critical patches on internet-facing systems. Firms that miss it usually own the right tools and lack the standing approval.

The two patching strategies also cover software the vendor no longer supports. At every level, online services the vendor no longer supports are removed, unsupported office suites, browsers, email clients, PDF software and security products are removed, and unsupported operating systems are replaced. Level 3 extends the removal to every unsupported application and asks that only the latest or the previous release of an operating system is in use.

What each level adds, strategy by strategy

The appendices list every requirement at each level, and the comparison in the model's Appendix D marks what each level adds in bold. Condensed to the point of each requirement, the additions are these.

STRATEGY	LEVEL 1	LEVEL 2 ADDS	LEVEL 3 ADDS
Patch applications	Asset discovery at least fortnightly. Online services scanned daily and the internet-content applications weekly. The fastest clock on critical flaws in online services; two weeks for the rest of the edge and for the internet-content applications. Unsupported online services and applications removed.	Every other application scanned at least fortnightly and patched within one month.	The fastest clock on critical flaws in the internet-content applications. Every unsupported application removed.
Patch operating systems	Internet-facing servers and network devices scanned daily, the rest fortnightly. The fastest clock on critical flaws at the edge, two weeks otherwise; workstations and internal devices within one month. Unsupported operating systems replaced.	No change from Level 1.	The fastest clock on critical flaws in workstations and internal devices. Drivers and firmware scanned fortnightly and patched on the same two clocks. Only the latest or the previous release of an operating system in use.
Multi-factor authentication	Required for users of the organisation's and third parties' online services that hold sensitive data, and where available for the rest. Required for customers of online services that hold their sensitive data. The factors are something the user has and knows, or has and unlocks by something known or inherent.	Required for privileged and unprivileged users of systems. Phishing-resistant for users of online services and of systems; a phishing-resistant option for customers. Authentication events centrally logged, protected and analysed; incidents reported and the response plan enacted.	Required, and phishing-resistant, for users of data repositories and for customers. Logs from internal servers and workstations analysed as well as the edge.
Restrict administrative privileges	Privileged access validated when first requested. Dedicated privileged accounts, kept off the internet, email and web unless explicitly authorised and then limited. Separate privileged and unprivileged operating environments, with no sign-in across the line.	Privileged access revalidated on a clock and removed when idle. Privileged environments never virtualised inside unprivileged ones. Administration through jump servers. Long, unique, managed credentials for break glass, local administrator and service accounts. Privileged access and group changes centrally logged.	Privileged access limited to what the duty requires. Secure administrative workstations. Just-in-time administration. The operating system's memory integrity and credential protection features enabled. Internal server and workstation logs analysed.

STRATEGY	LEVEL 1	LEVEL 2 ADDS	LEVEL 3 ADDS
Application control	On workstations, applied to user profiles and the temporary folders of the operating system, browsers and email clients. Executables, libraries, scripts, installers, compiled HTML, HTML applications and control panel applets restricted to an approved set.	On internet-facing servers and in every other location. Microsoft's recommended application blocklist. Rulesets validated at least annually. Allowed and blocked events centrally logged.	On internal servers. Drivers restricted to an approved set, with Microsoft's vulnerable driver blocklist. Internal server and workstation logs analysed.
Restrict Microsoft Office macros	Disabled for users without a demonstrated business need. Blocked in files from the internet. Scanned by antivirus. Settings users cannot change.	Macros blocked from making Win32 API calls.	Only macros from a sandbox, a Trusted Location or a trusted publisher's signature run. Macros checked for malicious code before signing or placement, and only the checkers can write to Trusted Locations. Untrusted or older signatures cannot be enabled from the Message Bar or Backstage View. The trusted publisher list validated at least annually.
User application hardening	The legacy browser disabled or removed. Browsers do not process Java or advertisements from the internet, and users cannot change browser security settings.	Browsers, office suites and PDF software hardened to ASD and vendor guidance, the most restrictive winning, with settings users cannot change. Office blocked from child processes, executable content, code injection and embedded package activation; PDF software blocked from child processes. Script and command line events centrally logged.	The old application framework and the old shell engine disabled or removed. The shell in constrained language mode. Internal server and workstation logs analysed.
Regular backups	Data, applications and settings backed up and retained to business need, synchronised to a common point in time, held securely and resiliently, and restored in disaster recovery exercises. Unprivileged accounts cannot reach other users' backups or modify or delete any backup.	Privileged accounts, except backup administrators, cannot reach other users' backups or modify or delete any backup.	No account can reach its own backups except a backup administrator, and backup administrators cannot modify or delete a backup during its retention period.

Condensed from Appendices A to D of the November 2023 edition. Each level includes everything in the level below it.

From Level 2, privileged access is disabled after 12 months unless it is revalidated, and after 45 days of inactivity. The two clocks turn a privileged user list into a living document: an access review that runs on the calendar, and an

idle account that switches itself off.

From Level 2 every strategy except backups carries the same logging tail: the strategy's events are centrally logged, the logs are protected from modification and deletion, events from internet-facing servers are analysed in a timely manner, incidents are reported to the chief information security officer or a delegate and to ASD as soon as possible, and the incident response plan is enacted. Level 3 extends the analysis to internal servers and workstations. From Level 2 a control with no log of its own events does not meet the model, whatever the setting says.

Choosing the level to aim at

The model asks an organisation to pick its target by how desirable it is to an attacker and what a compromise would cost. For most small firms the target is Level 2. The attacker Level 1 describes is real and common, but the attacker Level 2 describes is the one behind business email compromise, the phished multi-factor prompt and the deleted backup, and those are the incidents that close small firms. Level 2 is also the first level whose requirements a buyer can verify from logs, which is why questionnaires written by government and enterprise procurement teams tend to describe it.

Level 3 is the right target for a firm that holds health, financial or legal records in volume, that builds or hosts software for others, or that supplies government. The additions are affordable for a firm that already holds Level 2: just-in-time administration and secure administrative workstations are configuration on identities the firm already separates, driver control and the fastest clock on workstations are an extension of controls already running, and the extra log analysis is the same pipeline pointed at more sources.

Two rules hold whatever the target. The firm holds the level of its lowest strategy, so the plan takes every strategy up together. And every requirement is read back from the system that enforces it, so a setting the firm cannot export from that system is not counted, whatever its policy document says.

> The level a firm holds is the level of its lowest strategy, read from the systems that enforce the controls, on the day the assessor asks.

What to do

For a firm working toward a target level, the order that costs least and stops most:

1. Rate every strategy at its current level from the enforcing system, with the exceptions listed and dated, and fix the lowest first.
2. Put phishing-resistant multi-factor authentication on every account, starting with administrators and email, and block the sign-ins that bypass it.
3. Lock the backup copy so no ordinary or privileged account can delete it, and time a restore of real data.
4. Give every administrator a separate privileged account and a separate environment to use it from, then make privileged roles something they activate for a task.
5. Write the patch process down: who learns of a critical release, who holds the standing approval, how the change is applied inside the window on a weekend. Then set the scanner to the model's cadence.
6. Block macros from the internet, harden the browser, office suite and PDF reader to vendor guidance, and turn on central logging of script and command line events.

7. Run application control in audit, read what it would have blocked, admit the publishers you trust and the software you build, and enforce when the audit is quiet.
8. Give every control a check that runs every day and tells someone when its state changes.

We assess clients against the same edition: all eight strategies rated at the level they reach, read from the live systems, with a written report and an uplift plan sequenced in the order above and priced by strategy. The uplift is delivered in that order, and each control is read back at the end. Our free Essential Eight check at blackshard.com.au/tools/essential-eight-check gives a first reading in the browser, with nothing sent to us. Scope and price for the assessment are at blackshard.com.au/essential-eight-assessment.

Our record

Black Shard builds and defends the systems Australian organisations run on, and we assess clients against the November 2023 edition of the model.

The record is as held at 16 November 2026, the date this edition was revised.

- SMB1001:2026 Gold, issued by CyberCert, valid 16 Jun 2026 to 17 Jun 2027.
- Essential Eight Maturity Level 3, assessed September 2026.
- Prequalified supplier to Queensland councils under Local Buy arrangement LB308, ICT Solutions, Products, Services and New Technologies.
- Prequalified supplier under the NSW Government ICT Services Scheme (SCM0020).
- Holder of Victorian Fair Jobs Code pre-assessment certificate FJC-261023-11119, required for Victorian Government threshold and high-value procurement.
- Approved supplier on the Tasmanian Government Technology Services List (TSL) for Application Development Services, Operational Technical Services and Cyber Security.
- Our CISO holds the Exemplar Global ISO/IEC 27001:2022 Lead Auditor credential.
- Professional indemnity, public liability, cyber and workers' compensation insurance, all current.
- More than 65,000 tests across our repositories, counted on 26 October 2026.
- More than 1,500 deploys and releases shipped since 27 April 2026.

> 24/7 incident line: 1800 370 270.

Sources

- Australian Signals Directorate, Essential Eight maturity model (November 2023), published 27 November 2023: https://www.cyber.gov.au/business-government/asds-cyber-security-frameworks/essential-eight/essential-eight-maturity-model
- Australian Signals Directorate, Essential Eight maturity model changes (November 2023), published 27 November 2023: https://www.cyber.gov.au/business-government/asds-cyber-security-frameworks/essential-eight/essential-eight-maturity-model-changes