AEGIS EIGHT
aegiseight.com

ASD Essentials assessment

Automated Microsoft 365 posture review, aligned to the ASD Essentials maturity model.

Achieved maturity

ML0

FREE TRIAL — one report per Microsoft 365 tenant. Subscribe to enable continuous monitoring and change-triggered alerts.

Organisation	Contoso Pty Ltd
Contact	security@contoso.com.au
Tenant ID	contoso-tenant-id
Scan ID	sample-scan-001
Generated	26 June 2026 at 9:13 pm
Target maturity	ML2
Achieved maturity	ML0
Assessor	Aegis Eight automated assessment
Report version	1.0

Executive snapshot

Contoso Pty Ltd · scan 2026-06-26 · remediation roadmap on the next page · full evidence in section 4.

OVERALL MATURITY

Achieved **ML0** against a target of **ML2** · **Target not met**

Essentials uses a weakest-link rule: your overall level is set by the lowest-scoring strategy, not the average.

ASSESSMENT LIMITATIONS & EVIDENCE GAPS

This report is an automated Microsoft Graph posture review at the ASD **Fair** evidence tier — configuration reviewed via API responses, not interactive device testing. It is suitable for prioritising remediation but is not a substitute for a formal on-site or IRAP assessment.

Tier	What this means for you
Fair	Automated scan — live tenant configuration via Microsoft Graph. This report operates here.
Good	Applied-state verification, sign-in analysis, or RSoP exports — available via attestation or upcoming collectors (see "Upgrade: Good" hints in section 4).
Excellent	Simulated control testing (e.g. macro execution, application-block tests) — not automated in this pass.

Upgrade opportunity: opt-in synthetic tests can uplift MFA, admin isolation, and application control to **Excellent**-tier evidence — see the Synthetic testing section later in this report.

This scan has **5 no-visibility findings** and **5 poor-evidence findings**. Scope exclusions and next steps are listed in section 2.

- Evidence is collected from API responses (Fair tier per ASD process guide) — devices and policies were not interactively tested.
- Application control, Office macros, and user application hardening are assessed via Intune configuration profile heuristics (policy presence). Enforcement and per-device applied state require additional evidence where noted.
- Full scope and limitations in section 2.

Top risks

Business email compromise risk (3) · CRITICAL · BLOCKS ML1
Weak MFA coverage is the leading cause of account takeover and is a baseline requirement for most cyber-insurance policies.

Unpatched device exposure (9) · CRITICAL · BLOCKS ML1
Unpatched operating systems are a primary ransomware entry point and can affect cyber-insurance eligibility.

Data loss and recovery risk (2) · CRITICAL · BLOCKS ML1
Unverified backups mean a ransomware event could cause permanent data loss and extended downtime.

Single-account ransomware recovery failure (2) · CRITICAL
Unverified backups mean a ransomware event could cause permanent data loss and extended downtime.

Phishing-vulnerable authentication · HIGH
Weak MFA coverage is the leading cause of account takeover and is a baseline requirement for most cyber-insurance policies.

REMEDIATION ROADMAP

Path from ML0 to ML2: **0 of 10 completed · 6 in Immediate blockers to ML1 · baseline scan**. Future scans will show progress as actions are completed. Full checklist on the next page.

Remediation roadmap

> **0 of 10 completed · 6 in Immediate blockers to ML1 · baseline scan** · Mark actions complete in your runbook; the next scan will track progress.

A single prioritised path from ML0 to ML2. Complete ML-tier blockers first. Effort colours: **Quick fix** · **Multi-step** · **Day's work** · **Project**. Detailed recommendations are in section 4.

Immediate blockers to ML1

6 actions · Scale: Half day to one day

- [ ] **Protect missing Microsoft 365 workloads · Blocks ML1**
  Multi-step
- [ ] **Register MFA for 2 users · Blocks ML1**
  Quick fix
- [ ] **Remediate 3 non-compliant devices · Blocks ML1**
  Multi-step
- [ ] **Provide vulnerability scanning evidence · Blocks ML1**
  Day's work
- [ ] **Register MFA for affected users · Blocks ML1**
  Quick fix
- [ ] **Remediate affected non-compliant devices · Blocks ML1**
  Multi-step

Additional steps to ML2

4 actions · Scale: Major project

- [ ] **Separate backup admin from tenant admin**
  Quick fix
- [ ] **Deploy user application hardening baseline**
  Day's work
- [ ] **Enforce application control on workstations**
  Project
- [ ] **Separate Backup Administrator from Global Administrator**
  Quick fix

Cyber risk register

Unique risks grouped by theme (max 8 shown). Effort matches the roadmap legend above. Per-control evidence is in section 4.

Severity	Priority	Risk	Why it matters	Effort	ML	Count	Status
Critical	Immediate	**Business email compromise risk**	3 related findings — 2 users can sign in to Microsoft 365 with only a password, increasing the likelihood of…	Quick fix	**ML1**	3	Open
Critical	Immediate	**Unpatched device exposure**	9 related findings — 3 managed devices reports non-compliant and may be missing required security updates or…	Day's work	**ML1**	9	Open
Critical	Immediate	**Data loss and recovery risk**	2 related findings — One or more critical workloads (mail, SharePoint, OneDrive) lack demonstrated backup…	Multi-step	**ML1**	2	Open
Critical	Next 30 days	**Single-account ransomware recovery failure**	2 related findings — One account can both administer backups and hold Global Administrator, weakening the…	Quick fix	—	2	Open
High				Day's work	—	1	Open

	Next 30 days	**Phishing-vulnerable authentication**	SMS, voice, or email OTP defaults are easier to intercept than authenticator apps, FIDO2 keys, or Windows…				
Info	Monitor	**Standing privileged access risk**	Permanent privileged assignments mean attackers who compromise an admin account retain access until manually…	Project	—	1	Open
Info	Monitor	**Malicious software execution risk**	2 related findings — Untrusted executables may still run on endpoints — application control is a primary…	Project	—	2	Open
Info	Monitor	**Browser-borne malware exposure**	2 related findings — Browsers and productivity apps may lack ASD-aligned hardening — increasing success of…	Day's work	—	2	Open

22 open findings total · 8 unique risk groups.

Executive dashboard

Achieved **ML0** to target **ML2** · **Target not met**

Progress toward ML2

Multi-factor authentication	57% effective · achieved ML0	4/7
Restrict administrative privileges	100% effective · achieved ML2	4/4
Patch operating systems	33% effective · achieved ML0	2/6
Patch applications	40% effective · achieved ML0	2/5
Regular backups	50% effective · achieved ML0	3/6
Application control	67% effective · achieved ML1	2/3
Configure Microsoft Office macros	100% effective · achieved ML3	3/3
User application hardening	67% effective · achieved ML1	2/3

ASSESSMENT EVIDENCE

9 Microsoft Graph sources queried · 45 of 50 findings backed by live tenant data · Assessment performed 26 June 2026, 11:13 am (UTC). Full audit trail available on request via the auditor PDF.

Synthetic testing (Excellent tier) — not enabled

COVERAGE GAP — UPGRADE AVAILABLE

This assessment did not include live synthetic tests. ASD's Essentials process guide reserves the **Excellent** evidence tier for controls demonstrated by simulated activity. Aegis Eight can run three opt-in synthetic tests to uplift the evidence quality of your strongest controls:

- **MFA enforcement** — attempt sign-in as a sentinel account with no MFA, verify Conditional Access blocks or challenges it.
- **Admin internet isolation** — attempt sign-in as a privileged sentinel from the internet, verify CA blocks the session.
- **Application control** — verify the bundled Intune script reports an unsigned executable was blocked on a sentinel device.

To enable, follow **docs/synthetic-testing-runbook.md** (provision an **AegisEight-Synthetic** group, two cloud-only sentinel users, and a Windows Intune device), then enable synthetic testing from the Aegis Eight dashboard or the pre-scan setup wizard on your next trial.

1. Executive summary

Aegis Eight conducted an automated review of Contoso Pty Ltd's Microsoft 365 tenant against the ASD Essentials maturity model. The overall achieved maturity level is **ML0** against a target of **ML2**. The tenant does not currently meet the target maturity level based on automated signals.

> **WHY YOUR OVERALL MATURITY IS ML0**
> Essentials maturity follows the **weakest-link rule**: your overall level is set by the lowest-scoring mitigation strategy, not the average. Findings rated Effective elsewhere do not raise the overall number until every strategy reaches the next maturity level together.
>
> **Blocking ML1:** Multi-factor authentication, Patch operating systems, Patch applications, and Regular backups — overall maturity rises to ML1 once every one of these reaches at least ML1.

Mitigation strategy summary

Mitigation strategy	Achieved	Worst outcome
Multi-factor authentication	ML0	Ineffective
Restrict administrative privileges	ML2	Ineffective
Patch operating systems	ML0	Ineffective
Patch applications	ML0	Ineffective
Regular backups	ML0	Ineffective
Application control	ML1	Ineffective
Configure Microsoft Office macros	ML3	Effective
User application hardening	ML1	Ineffective

2. Scope

Automated Essentials posture review of a Microsoft 365 tenant via the Microsoft Graph API. The assessment is conducted using app-only or delegated OAuth credentials and reviews configuration accessible to those credentials.

In scope

• Microsoft 365 tenant identity, authentication methods, and directory roles via Microsoft Graph.
• Microsoft Intune managed devices and their reported compliance state.
• Privileged account assignments visible in Microsoft Entra ID.
• Microsoft 365 Backup protection policies, restore points, restore sessions, and Backup Administrator role assignments via Microsoft Graph (when the tenant has Microsoft 365 Backup enabled).

Out of scope (this automated pass)

• On-premises infrastructure, third-party identity providers, and customer-facing applications.
• Application control policies (AppLocker/WDAC), Office macro configuration, and user application hardening baselines — these require RSoP exports or Intune configuration profile review.

• Third-party backup tooling (Veeam, Datto, Acronis, AvePoint, etc.), backup immutability/off-tenant copies, and signed-off DR exercise records — provide attestation to elevate evidence quality.
• Vulnerability scanner cadence, patch timeframes against ASD windows, and event log forwarding/retention.

Assessment limitations

• Evidence is collected from API responses (Fair tier per ASD process guide) — devices and policies were not interactively tested.
• Application control, Office macros, and user application hardening are assessed via Intune configuration profile heuristics (policy presence). Enforcement and per-device applied state require additional evidence where noted.
• Backup posture is automated for Microsoft 365 Backup tenants only — third-party backup products (Veeam, Datto, Acronis, etc.) require attested evidence.
• The scan is point-in-time. Continuous monitoring requires a subscription plan.

> **PERMISSIONS USED & DATA HANDLING**
>
> Aegis Eight accessed your Microsoft 365 tenant under **read-only Microsoft Graph permissions** granted by your Global Administrator on the Microsoft consent screen. No write, create, update, or delete permissions were requested or granted.
>
> • **Scopes used:** User.Read.All, Directory.Read.All, Policy.Read.All, RoleManagement.Read.Directory, UserAuthenticationMethod.Read.All, DeviceManagementManagedDevices.Read.All, DeviceManagementConfiguration.Read.All, AuditLog.Read.All, BackupRestore-Configuration.Read.All, BackupRestore-Restore.Read.All.
>
> • **Tenant isolation:** access is scoped to your tenant only. Multi-tenancy means Aegis Eight is deployed into many tenants, not that it reads across them.
>
> • **Retention:** scan signals and raw API responses are retained for 90 days (full payload), then reduced to hashes and summary metadata. Per-tenant shorter-retention overrides are available on request.
>
> • **Revocation:** Entra admin centre › Enterprise applications › Aegis Eight › Properties › Delete. On request after revocation, all tenant scan data is purged within 30 days.
>
> Full details: aegiseight.com/security.

3. Methodology

This assessment follows the ASD Essentials assessment process guide (October 2024). Planning and scope are established when you connect the tenant; control assessment runs by querying Microsoft Graph for configuration evidence; this report is the development stage.

Automated collection operates at the ASD "Fair" evidence tier — reviewing system configuration through API responses rather than running scripts on individual devices. Findings in the mitigation strategies section include a plain-language "What this means" line for gaps that need attention.

Each sub-control receives one of seven standardised outcomes. A mitigation strategy is only claimed at a maturity level when every sub-control at that level is rated Effective or Alternate control.

4. Mitigation strategies

The eight ASD mitigation strategies are presented below. Each strategy includes sub-control findings with their assessment outcome, evidence quality, and observation. Affected account and device lists are in the appendix.

4.1 Multi-factor authentication

Achieved: **ML0** · Evidence: **Good** · 4/8 controls effective.

ML	Control	Outcome	Evidence	Observation
ML1	Multi-factor authentication is used to authenticate privileged users of important data repositories and online services.	Effective	Fair	All 3 privileged user(s) have MFA registered.
ML1	Multi-factor authentication is used to authenticate the organisation's users of online services that process, store or communicate sensitive data.	Ineffective	Fair	26/28 users (92.9%) have MFA registered against the Microsoft Entra tenant. *What this means: 2 users can sign in to Microsoft 365 with only a password, increasing the likelihood of account takeover, email fraud, and access to SharePoint, Teams, and OneDrive data.*
ML1	Multi-factor authentication is used to authenticate users of online customer services that process, store or communicate sensitive customer data.	No visibility	Poor	Customer-facing services and any non-M365 identity providers are outside the visibility of this automated scan. *What this means: Can sign in to Microsoft 365 with only a password, increasing the likelihood of account takeover, email fraud, and access to SharePoint, Teams, and OneDrive data.*
ML2	Multi-factor authentication is used to authenticate unprivileged users of important data repositories and internal systems.	Ineffective	Fair	Registration: 92.9%, capability coverage: 100.0%. *What this means: 2 users can sign in to Microsoft 365 with only a password, increasing the likelihood of account takeover, email fraud, and access to SharePoint, Teams, and OneDrive data.*
ML2	Successful and unsuccessful multi-factor authentication events are centrally logged.	Effective	Good	412 sign-in event(s) retrieved from Microsoft Entra over the last 7 day(s) (397 successful, 287 required MFA).
ML3	Multi-factor authentication is phishing-resistant.	Ineffective	Fair	67.9% of users have a phishing-resistant default method (authenticator, FIDO2, Windows Hello, passkey). *What this means: SMS, voice, or email OTP defaults are easier to intercept than authenticator apps, FIDO2 keys, or Windows Hello — weakening protection against targeted phishing.*
ML1	Privileged users are actually challenged for MFA on real sign-ins (not just registered for MFA).	Effective	Good	22/22 privileged sign-ins required MFA over the last 7 day(s); no successful privileged sign-ins bypassed MFA.

ML1	Conditional Access requires MFA for organisation user sign-ins (Microsoft what-if simulation).	Effective	Good	What-if simulation for user1@contoso.com.au: 2 enabled policy/policies would apply; MFA enforced by: Require MFA for all users.

Affected accounts and devices

• [ML1] Multi-factor authentication is used to authenticate the organisation's users of… — 2 affected
• [ML2] Multi-factor authentication is used to authenticate unprivileged users of… — 2 affected
• [ML3] Multi-factor authentication is phishing-resistant. — 9 affected
Full lists are in the appendix at the end of this report.

Recommendations

• **[ML1]** Drive MFA registration above 95% — 2 user(s) still unregistered.

• **[ML1]** Provide evidence of MFA enforcement on any customer-facing application that handles sensitive data.

• **[ML2]** Push registration to at least 99% and ensure all users are MFA-capable.

• **[ML2]** Confirm Microsoft Entra sign-in logs are forwarded to a centralised log store (e.g. Sentinel, Splunk) with retention meeting ASD ML2 requirements.

• **[ML3]** Migrate users from SMS/voice/email defaults to authenticator app, FIDO2 security keys, or Windows Hello for Business.

4.2 Restrict administrative privileges

Achieved: **ML2** · Evidence: **Good** · 4/5 controls effective.

ML	Control	Outcome	Evidence	Observation
ML1	Privileged accounts are strictly limited to only those required to administer the system.	Effective	Fair	3 effective Global Administrator(s) (3 direct user assignment(s), 0 group-based privileged assignment(s)); 6 total privileged role assignment(s) in Microsoft Entra.
ML1	Privileged accounts are prevented from accessing the internet, email and web services from untrusted networks (simulated via Conditional Access).	Effective	Good	What-if for user0@contoso.com.au from simulated untrusted location: access would be blocked by Block privileged roles from non-trusted locations.
ML2	Privileged access is limited to that required for personnel to undertake their duties (ML2).	Effective	Fair	3 effective Global Administrator(s) (3 direct); ML2 expects no more than 3.
ML2	Privileged role activations are recorded (PIM assignment schedule requests).	Effective	Good	12 privileged PIM activation(s) recorded in the last 90 days.

ML3	Just-in-time administration is used — privileged roles are eligible for activation, not permanently assigned.	Ineffective	Good	3 permanent Global Administrator(s) via directory roles. 3 eligible (PIM) Global Administrator assignment(s). 1 time-bound active Global Administrator activation(s) at scan time. 12 privileged activation request(s) in the last 90 days. *What this means: Permanent privileged assignments mean attackers who compromise an admin account retain access until manually revoked.*

Recommendations

• **[ML2]** Ensure activation logs are forwarded to your SIEM with privileged-access monitoring.

• **[ML3]** Convert standing Global Administrator assignments to eligible PIM assignments with approval and time-bound activation.

4.3 Patch operating systems

Achieved: **ML0** · Evidence: Good · 2/7 controls effective.

ML	Control	Outcome	Evidence	Observation
ML1	Missing-patch risk is tracked with vulnerability scanning (not inferred from Intune compliance alone).	No visibility	Poor	This scan does not collect CVE/KB age or scanner cadence — Intune compliance is a separate proxy finding below. *What this means: Intune compliance is a proxy for patch state, not direct proof of missing KBs or CVEs. Provide Defender Vulnerability Management or equivalent scanner output to demonstrate patch cadence against ASD windows.*
ML1	Operating systems patch timeliness is assessed via Intune compliance as a proxy (not direct CVE/KB age).	Ineffective	Fair	Intune compliance used as a proxy for ASD patch timeframes — 19/22 devices (86.4%) report compliant against Intune policy; 3 non-compliant. *What this means: 3 managed devices reports non-compliant and may be missing required security updates or configuration — confirm patch state with vulnerability management evidence.*
ML2	ML2 patch timeframes are assessed via Intune compliance proxy (not direct vulnerability age).	Ineffective	Fair	19/22 devices (86.4%) report compliant against Intune policy; 3 non-compliant. *What this means: 3 managed devices reports non-compliant and may be missing required security updates or configuration — confirm patch state with vulnerability management evidence.*

ML3	ML3 patch timeframes are assessed via Intune compliance proxy (not direct vulnerability age).	Ineffective	Fair	19/22 devices (86.4%) report compliant against Intune policy; 3 non-compliant. *What this means: 3 managed devices reports non-compliant and may be missing required security updates or configuration — confirm patch state with vulnerability management evidence.*
ML1	Operating systems patch status verified via Intune compliance-policy RSoP (not only device-reported complianceState).	Effective	Good	26/28 sampled device(s) report compliant RSoP for all assigned compliance policies.
ML2	Operating systems meet ML2 patch timeframes with RSoP-verified compliance.	Effective	Good	26/28 sampled device(s) report compliant RSoP for all assigned compliance policies.
ML1	Windows Update for Business or equivalent OS update policies are deployed.	Ineffective	Fair	No Windows Update for Business or patch-OS Intune policies were detected. *What this means: Reports non-compliant and may be missing required security updates or configuration — confirm patch state with vulnerability management evidence.*

Affected accounts and devices

• [ML1] Operating systems patch timeliness is assessed via Intune compliance as a proxy… — 3 affected
• [ML2] ML2 patch timeframes are assessed via Intune compliance proxy (not direct… — 3 affected
• [ML3] ML3 patch timeframes are assessed via Intune compliance proxy (not direct… — 3 affected

Full lists are in the appendix at the end of this report.

Recommendations

• **[ML1]** Grant ThreatHunting.Read.All and enable Defender TVM, or provide third-party vulnerability evidence.

• **[ML1]** Investigate 3 non-compliant device(s); compliance gaps are a proxy signal, not proof of missing patches.

• **[ML2]** Raise device compliance to at least 95% for ML2.

• **[ML3]** Raise device compliance to at least 99% for ML3.

• **[ML1]** Deploy WUfB update rings via Intune for OS patch management.

4.4 Patch applications

Achieved: **ML0** · Evidence: Good · 2/6 controls effective.

ML	Control	Outcome	Evidence	Observation
ML1	Missing-patch risk is tracked with vulnerability scanning (not inferred from Intune compliance alone).	No visibility	Poor	This scan does not collect CVE/KB age or scanner cadence — Intune compliance is a separate proxy finding below. *What this means: Intune compliance is a proxy for patch state, not direct proof of missing KBs or CVEs. Provide Defender Vulnerability Management or equivalent scanner output to demonstrate patch cadence against ASD windows.*

ML1	Applications patch timeliness is assessed via Intune compliance as a proxy (not direct CVE/KB age).	Ineffective	Fair	Intune compliance used as a proxy for ASD patch timeframes — 19/22 devices (86.4%) report compliant against Intune policy; 3 non-compliant. *What this means: 3 managed devices reports non-compliant and may be missing required security updates or configuration — confirm patch state with vulnerability management evidence.*
ML2	ML2 patch timeframes are assessed via Intune compliance proxy (not direct vulnerability age).	Ineffective	Fair	19/22 devices (86.4%) report compliant against Intune policy; 3 non-compliant. *What this means: 3 managed devices reports non-compliant and may be missing required security updates or configuration — confirm patch state with vulnerability management evidence.*
ML3	ML3 patch timeframes are assessed via Intune compliance proxy (not direct vulnerability age).	Ineffective	Fair	19/22 devices (86.4%) report compliant against Intune policy; 3 non-compliant. *What this means: 3 managed devices reports non-compliant and may be missing required security updates or configuration — confirm patch state with vulnerability management evidence.*
ML1	Applications patch status verified via Intune compliance-policy RSoP (not only device-reported complianceState).	Effective	Good	26/28 sampled device(s) report compliant RSoP for all assigned compliance policies.
ML2	Applications meet ML2 patch timeframes with RSoP-verified compliance.	Effective	Good	26/28 sampled device(s) report compliant RSoP for all assigned compliance policies.

Affected accounts and devices

• [ML1] Applications patch timeliness is assessed via Intune compliance as a proxy (not… — 3 affected
• [ML2] ML2 patch timeframes are assessed via Intune compliance proxy (not direct… — 3 affected
• [ML3] ML3 patch timeframes are assessed via Intune compliance proxy (not direct… — 3 affected
Full lists are in the appendix at the end of this report.

Recommendations

• **[ML1]** Grant ThreatHunting.Read.All and enable Defender TVM, or provide third-party vulnerability evidence.
• **[ML1]** Investigate 3 non-compliant device(s); compliance gaps are a proxy signal, not proof of missing patches.
• **[ML2]** Raise device compliance to at least 95% for ML2.
• **[ML3]** Raise device compliance to at least 99% for ML3.

4.5 Regular backups

Achieved: **ML0** · Evidence: Good · 5/9 controls effective.

ML	Control	Outcome	Evidence	Observation

ML1	Backups of important data, software and configuration settings are performed and retained.	Ineffective	Fair	2 Microsoft 365 Backup retention policy(ies) covering exchange, oneDriveForBusiness; 56 protected item(s). *What this means: One or more critical workloads (mail, SharePoint, OneDrive) lack demonstrated backup coverage — increasing recovery time after ransomware or accidental deletion.*
ML1	Recovery points exist so backed-up data can be restored.	Effective	Fair	412 restore point(s) observed across sampled protection units.
ML1	Restoration is tested as part of disaster recovery exercises.	Effective	Fair	1 restore session(s) recorded in the last 365 days; most recent 2026-04-12T03:24:00Z.
ML1	A dedicated Backup Administrator role isolates backup access from other privileged duties.	Effective	Fair	1 Backup Administrator assignment(s).
ML2	Privileged accounts (excluding backup administrators) cannot modify or delete backups.	Ineffective	Fair	1 account(s) hold both Global Administrator and Backup Administrator roles, so a single compromise grants both backup write and tenant-wide admin. *What this means: One account can both administer backups and hold Global Administrator, weakening the principle that recovery credentials survive a tenant-wide compromise.*
ML2	Unprivileged accounts cannot access backups belonging to other accounts and cannot modify or delete backups.	No visibility	Poor	Unprivileged-account access to the M365 Backup admin centre / vault is not visible through Graph alone — vault RBAC and tenant settings must be reviewed manually. *What this means: Holds both Global Administrator and Backup Administrator roles, so one compromised credential could both control the tenant and weaken backup recovery — a common ransomware pattern.*
ML3	Privileged Backup Administrator access is limited and just-in-time, with no permanent standing assignments beyond the minimum needed.	Effective	Fair	1 permanent Backup Administrator assignment(s); ML3 expects no more than 2 and prefers PIM-eligible (not permanent) assignments.
ML3	Backups are protected against destructive synchronisation, ransomware, and tenant-level compromise (immutability / off-site copy).	No visibility	Poor	Microsoft 365 Backup is logically isolated from primary workload deletion, but tenant-level immutability, off-tenant copies, and ransomware-recovery posture cannot be confirmed via Graph. *What this means: One or more critical workloads (mail, SharePoint, OneDrive) lack demonstrated backup coverage — increasing recovery time after ransomware or accidental deletion.*
ML3	Backup Administrator privilege uses just-in-time activation where possible.	Effective	Good	1 permanent Backup Administrator(s) via directory roles; 1 PIM-eligible; 0 active JIT.

Affected accounts and devices

• [ML1] Backups of important data, software and configuration settings are performed… — 1 affected
• [ML2] Privileged accounts (excluding backup administrators) cannot modify or delete… — 1 affected
Full lists are in the appendix at the end of this report.

Recommendations

• **[ML1]** Configure Microsoft 365 Backup protection policies for: SharePoint (sites).

• **[ML2]** Remove Backup Administrator from accounts that also hold Global Administrator. Use dedicated backup-only accounts.

• **[ML2]** Confirm that only Backup Administrators (and not standard users) can access the Microsoft 365 Backup admin centre or the third-party backup console.

• **[ML3]** Confirm Microsoft Entra PIM is configured for time-bound, approval-based elevation of the Backup Administrator role.

• **[ML3]** Provide evidence of immutability / off-tenant retention (Microsoft documentation, third-party vendor configuration export, or signed attestation).

4.6 Application control

Achieved: **ML1** · Evidence: **Good** · Intune configuration profiles detected for application control (presence only).

ML	Control	Outcome	Evidence	Observation
ML1	Application control is implemented on workstations and servers per ASD maturity model.	Effective	Fair	1 Intune profile(s) matched application control heuristics: WDAC — allow signed catalog apps.
ML2	Application control blocks execution of executables unless they meet ASD allow-list criteria.	Effective	Fair	Policy presence is visible via Graph; audit vs enforced mode and per-device applied state are not verified in v1. 1 matching profile(s) present.
ML3	Application control rulesets are validated and maintained per ASD maturity model.	Effective	Fair	Policy presence is visible via Graph; audit vs enforced mode and per-device applied state are not verified in v1. ML3 ruleset depth (allow-lists, path rules) cannot be validated from profile metadata alone.
ML2	Application control blocks execution of executables unless they meet ASD allow-list criteria. (verified via per-device RSoP)	Ineffective	Good	25/28 sampled device(s) report succeeded RSoP for application control profile(s) (28 device(s)). *What this means: Untrusted executables may still run on endpoints — application control is a primary defence against commodity malware and ransomware droppers.*

ML3	Application control rulesets are validated and maintained per ASD maturity model. (verified via per-device RSoP)	Ineffective	Good	25/28 sampled device(s) report succeeded RSoP for application control profile(s) (28 device(s)). Ruleset depth still requires policy export for ML3 attestation. *What this means: Untrusted executables may still run on endpoints — application control is a primary defence against commodity malware and ransomware droppers.*

Affected accounts and devices

• [ML1] Application control is implemented on workstations and servers per ASD maturity… — 1 affected
• [ML2] Application control blocks execution of executables unless they meet ASD… — 1 affected
• [ML3] Application control rulesets are validated and maintained per ASD maturity… — 1 affected
Full lists are in the appendix at the end of this report.

Recommendations

• **[ML1]** Review assigned profiles and confirm they are deployed to all in-scope application control devices.

• **[ML2]** Provide RSoP or deviceConfigurationStates export to prove policies are applied and enforced on workstations.

• **[ML3]** Export the full ruleset (AppLocker/WDAC XML or Intune settings catalog) and a recent compliance report.

• **[ML2]** Investigate devices where application control profiles did not reach succeeded state in Intune.

• **[ML3]** Raise succeeded RSoP coverage to 90%+ across the fleet.

4.7 Configure Microsoft Office macros

Achieved: **ML3** · Evidence: Good · Office macro settings: Intune profiles present and RSoP succeeded on sampled devices (Good-tier evidence).

ML	Control	Outcome	Evidence	Observation
ML1	Microsoft Office macros are configured per ASD maturity model (disabled or strictly controlled).	Effective	Fair	1 Intune profile(s) matched Office macro settings heuristics: Block Office macros from internet.
ML2	Macros from the internet are blocked; only macros from trusted locations may run.	Effective	Fair	Policy presence is visible via Graph; audit vs enforced mode and per-device applied state are not verified in v1. 1 matching profile(s) present.
ML3	Macro settings are centrally managed and verified across all Office applications.	Effective	Fair	Policy presence is visible via Graph; audit vs enforced mode and per-device applied state are not verified in v1. ML3 ruleset depth (allow-lists, path rules) cannot be validated from profile metadata alone.
ML2	Macros from the internet are blocked; only macros from trusted locations may run. (verified via per-device RSoP)	Effective	Good	28/28 sampled device(s) report succeeded RSoP for Office macro settings profile(s) (28 device(s)).

ML3	Macro settings are centrally managed and verified across all Office applications. (verified via per-device RSoP)	Effective	Good	28/28 sampled device(s) report succeeded RSoP for Office macro settings profile(s) (28 device(s)). Ruleset depth still requires policy export for ML3 attestation.

Affected accounts and devices

• [ML1] Microsoft Office macros are configured per ASD maturity model (disabled or… — 1 affected
• [ML2] Macros from the internet are blocked; only macros from trusted locations may… — 1 affected
• [ML3] Macro settings are centrally managed and verified across all Office… — 1 affected
Full lists are in the appendix at the end of this report.

Recommendations

• **[ML1]** Review assigned profiles and confirm they are deployed to all in-scope Office macro settings devices.

• **[ML2]** Provide RSoP or deviceConfigurationStates export to prove policies are applied and enforced on workstations.

• **[ML3]** Export the full ruleset (AppLocker/WDAC XML or Intune settings catalog) and a recent compliance report.

• **[ML3]** Export full ruleset (AppLocker/WDAC XML or settings catalog) for assessor review.

4.8 User application hardening

Achieved: **ML1** · Evidence: Good · Intune configuration profiles detected for user application hardening (presence only).

ML	Control	Outcome	Evidence	Observation
ML1	User applications (browser, PDF, Office, .NET, PowerShell) are hardened per ASD maturity model.	Effective	Fair	1 Intune profile(s) matched user application hardening heuristics: Microsoft Edge hardening baseline.
ML2	Hardening settings block legacy frameworks, risky browser features, and unapproved scripts.	Effective	Fair	Policy presence is visible via Graph; audit vs enforced mode and per-device applied state are not verified in v1. 1 matching profile(s) present.
ML3	Hardening baselines are validated and maintained across all in-scope applications.	Effective	Fair	Policy presence is visible via Graph; audit vs enforced mode and per-device applied state are not verified in v1. ML3 ruleset depth (allow-lists, path rules) cannot be validated from profile metadata alone.
ML2	Hardening settings block legacy frameworks, risky browser features, and unapproved scripts. (verified via per-device RSoP)	Ineffective	Good	24/28 sampled device(s) report succeeded RSoP for user application hardening profile(s) (28 device(s)). *What this means: Browsers and productivity apps may lack ASD-aligned hardening — increasing success of drive-by downloads and malicious links.*

ML3	Hardening baselines are validated and maintained across all in-scope applications. (verified via per-device RSoP)	Ineffective	Good	24/28 sampled device(s) report succeeded RSoP for user application hardening profile(s) (28 device(s)). Ruleset depth still requires policy export for ML3 attestation. *What this means: Browsers and productivity apps may lack ASD-aligned hardening — increasing success of drive-by downloads and malicious links.*

Affected accounts and devices

• [ML1] User applications (browser, PDF, Office, .NET, PowerShell) are hardened per ASD… — 1 affected
• [ML2] Hardening settings block legacy frameworks, risky browser features, and… — 1 affected
• [ML3] Hardening baselines are validated and maintained across all in-scope… — 1 affected
Full lists are in the appendix at the end of this report.

Recommendations

• **[ML1]** Review assigned profiles and confirm they are deployed to all in-scope user application hardening devices.

• **[ML2]** Provide RSoP or deviceConfigurationStates export to prove policies are applied and enforced on workstations.

• **[ML3]** Export the full ruleset (AppLocker/WDAC XML or Intune settings catalog) and a recent compliance report.

• **[ML2]** Investigate devices where user application hardening profiles did not reach succeeded state in Intune.

• **[ML3]** Raise succeeded RSoP coverage to 90%+ across the fleet.

Appendix — affected accounts and devices

42 affected accounts or devices across 20 findings. Summary by control:

- Multi-factor authentication — 13 affected
- Patch applications — 9 affected
- Patch operating systems — 9 affected
- Application control — 3 affected
- Configure Microsoft Office macros — 3 affected
- User application hardening — 3 affected
- Regular backups — 2 affected

Detail — highest-volume findings

MULTI-FACTOR AUTHENTICATION · [ML3] MULTI-FACTOR AUTHENTICATION IS PHISHING-RESISTANT. — 9 AFFECTED

• user10@contoso.com.au (sms)
• user13@contoso.com.au (sms)
• user16@contoso.com.au (sms)
• user19@contoso.com.au (sms)
• user1@contoso.com.au (sms)
• user22@contoso.com.au (sms)
• user25@contoso.com.au (sms)
• user4@contoso.com.au (sms)
• user7@contoso.com.au (sms)

PATCH OPERATING SYSTEMS · [ML1] OPERATING SYSTEMS PATCH TIMELINESS IS ASSESSED VIA INTUNE COMPLIANCE AS A PROXY (NOT DIRECT CVE/KB… — 3 AFFECTED

- IPAD-RECEPTION · iOS 17.4 [inGracePeriod]
- LAPTOP-FINANCE-01 · Windows 10.0.22631 · user12@contoso.com.au [noncompliant]
- MAC-DESIGN-04 · macOS 14.4 · user24@contoso.com.au [noncompliant]

PATCH OPERATING SYSTEMS · [ML2] ML2 PATCH TIMEFRAMES ARE ASSESSED VIA INTUNE COMPLIANCE PROXY (NOT DIRECT VULNERABILITY AGE). — 3 AFFECTED

- IPAD-RECEPTION · iOS 17.4 [inGracePeriod]
- LAPTOP-FINANCE-01 · Windows 10.0.22631 · user12@contoso.com.au [noncompliant]
- MAC-DESIGN-04 · macOS 14.4 · user24@contoso.com.au [noncompliant]

PATCH OPERATING SYSTEMS · [ML3] ML3 PATCH TIMEFRAMES ARE ASSESSED VIA INTUNE COMPLIANCE PROXY (NOT DIRECT VULNERABILITY AGE). — 3 AFFECTED

- IPAD-RECEPTION · iOS 17.4 [inGracePeriod]
- LAPTOP-FINANCE-01 · Windows 10.0.22631 · user12@contoso.com.au [noncompliant]
- MAC-DESIGN-04 · macOS 14.4 · user24@contoso.com.au [noncompliant]

PATCH APPLICATIONS · [ML1] APPLICATIONS PATCH TIMELINESS IS ASSESSED VIA INTUNE COMPLIANCE AS A PROXY (NOT DIRECT CVE/KB AGE). — 3 AFFECTED

- IPAD-RECEPTION · iOS 17.4 [inGracePeriod]
- LAPTOP-FINANCE-01 · Windows 10.0.22631 · user12@contoso.com.au [noncompliant]
- MAC-DESIGN-04 · macOS 14.4 · user24@contoso.com.au [noncompliant]

Other affected findings (counts only)

- [ML2] ML2 patch timeframes are assessed via Intune compliance proxy (not direct… — 3 affected
- [ML3] ML3 patch timeframes are assessed via Intune compliance proxy (not direct… — 3 affected
- [ML1] Multi-factor authentication is used to authenticate the organisation's users of… — 2 affected
- [ML2] Multi-factor authentication is used to authenticate unprivileged users of… — 2 affected
- [ML1] Backups of important data, software and configuration settings are performed… — 1 affected
- [ML2] Privileged accounts (excluding backup administrators) cannot modify or delete… — 1 affected
- [ML1] Application control is implemented on workstations and servers per ASD maturity… — 1 affected
- [ML2] Application control blocks execution of executables unless they meet ASD… — 1 affected
- [ML3] Application control rulesets are validated and maintained per ASD maturity… — 1 affected
- [ML1] Microsoft Office macros are configured per ASD maturity model (disabled or… — 1 affected
- [ML2] Macros from the internet are blocked; only macros from trusted locations may… — 1 affected
- [ML3] Macro settings are centrally managed and verified across all Office… — 1 affected
- [ML1] User applications (browser, PDF, Office, .NET, PowerShell) are hardened per ASD… — 1 affected
- [ML2] Hardening settings block legacy frameworks, risky browser features, and… — 1 affected
- [ML3] Hardening baselines are validated and maintained across all in-scope… — 1 affected

7. Conclusion and next steps

This automated assessment is intended to accelerate — not replace — a formal ASD Essentials assessment. Outcomes marked "No visibility" or "Not assessed" require manual evidence collection per the ASD process guide.

THIS IS YOUR BASELINE SCAN — WHAT UNLOCKS NEXT

Continuous-monitoring subscribers receive:

- **Change-triggered alerts** whenever an assessed control regresses — new admin without MFA, a device falling out of compliance, a backup policy disabled — so posture drift surfaces in days, not at the next quarterly audit.
- **Trend reporting** across re-scans, so future versions of this report show progress (or regression) over time rather than a single point-in-time snapshot.
- **Path 2 collectors** as they ship (sign-in log analysis, Intune applied-state, Conditional Access What-If, PIM activation history) — each one lifts the findings marked "Upgrade: Good" above from Fair to Good evidence quality, which materially strengthens IRAP-QAF and cyber-insurance evidence trails.
- **Evidence-chain retention** for the full subscription window (vs. 90 days on the free trial), with hash-addressed raw Microsoft Graph responses available on demand for assessors.

Subscribe at aegiseight.com to enable.

What Aegis Eight provides

Posture visibility · Automated assessment support · Evidence collection · Remediation guidance.

What Aegis Eight does not provide

Formal ASD certification · Guaranteed compliance · Legal sign-off. This report should be used in conjunction with a qualified assessor for any compliance attestation.

References

- ASD Essentials assessment process guide (October 2024)
- ASD Essentials maturity model (November 2023)