

YOUR MISSION

Move beyond assessment backlogs. Lead with clearer assurance.

A practical path for Australian Government CISOs to prepare evidence, focus assessor effort and support informed authorisation decisions.

About this case study

An anonymised composite of insights gathered while co-creating business cases across different clients. It is not a verified single-client success story. Figures and the logo's timing strapline are illustrative, not audited results or guarantees; readiness means drafts for expert review.

When assurance cannot keep pace

You are accountable for protecting government services while systems, suppliers and risks keep changing. Yet assessment backlogs and scarce assessor capacity can leave your team working from yesterday's evidence.

System Security Plans (SSPs), Statements of Applicability (SoAs), Cybersecurity Assessment Reports (CARs) and Security Risk Management Plans (SRMPs) can become outdated before the next review begins. Teams repeat document preparation, reconcile conflicting versions and chase evidence while authorisation decisions wait.

The pressure is operational and personal: explain the risk, support service delivery and give executives a defensible basis for decisions. Your team needs time for judgement, not just document production.

The cost of standing still

Without a current view of evidence and control effectiveness, risk drift can go unnoticed, remediation priorities can lose focus and authorisation delays can persist. Faster paperwork alone does not resolve these problems; reviewed evidence and accountable decisions do.

A GUIDE AND A CLEAR PLAN

Let AI assist. Keep people accountable.

DyDox is positioned as a Sovereign AI approach to assessment preparation and security documentation. It helps teams structure evidence and produce draft assessment and risk outputs, with human review at the centre of the workflow.

01

Agree the boundary

Define the system, assessment scope, evidence sources and success criteria. Confirm hosting, data residency, access and support arrangements. Where required, scope an optional offline or air-gapped deployment.

02

Prepare and assess

Use AI-assisted processing to organise supplied evidence, support control assessment and draft SSP, SoA, CAR and SRMP content. Reviewers check provenance, gaps, assumptions and proposed findings.

03

Review, decide and refresh

Assessors validate outputs and resolve exceptions. Risk owners agree treatment priorities; the authorising authority makes the decision. Refresh the evidence and dashboard as the environment changes.

What this could enable for your team

Capacity: redirect time from repeated preparation towards difficult findings and remediation.

Decision quality: use more consistent, reviewable documentation to explain risk and priorities.

Continuity: refresh the assurance picture as new evidence becomes available.

These are intended benefits, not observed outcomes from a verified deployment. Sovereignty and isolation depend on the agreed architecture and operating controls. AI assistance does not confer compliance, certification or authorisation.

MEDIUM-COMPLEXITY SCENARIO

Around 95% less illustrative effort

380 hours manual versus approximately 20 hours DyDox-assisted.
A planning illustration, not a measured productivity or delivery commitment.

Activity	Manual baseline	DyDox-assisted
Readiness	4 weeks / 160 h	2 h processing
Assessment	3 weeks / 120 h	1 h processing
SRMP	2 weeks / 80 h	1 h processing
Dashboard	0.5 week / 20 h	Real-time*
Human review	Within baseline**	2 days / 16 h
Illustrative total	9.5 weeks / 380 h	Approximately 20 h

* Real-time describes dashboard updates when evidence is available and processed, not continuous validation of every control. No separate dashboard effort is allocated in this illustration.

** Review is assumed to be included within the manual stage estimates; it is shown separately for the assisted scenario.

HOW TO READ THE NUMBERS

The manual baseline assumes 40-hour weeks; human review assumes two 8-hour days. The calculation is $(380 - 20) / 380 = 94.7\%$, rounded to around 95%. The assisted total adds four hours of AI processing to 16 hours of human review. It mixes processing duration and human effort, so it is not a like-for-like measured labour-saving result.

Processing time is not total elapsed time. Evidence collection, deployment, configuration, onboarding, integration, remediation, additional review, scheduling and authorisation may add time and are not separately modelled. The figures do not establish a 2.5-day end-to-end assessment or authorisation timeline.

Methodology: user-approved scenario assumptions informed by composite business-case discussions; no client-level dataset, sample size, controlled benchmark or independent audit is provided. Results depend on scope, evidence quality, model and hardware, deployment constraints and review depth. AI outputs may be incomplete or incorrect and require qualified validation.

Build a business case for your environment

[Start your DyDox journey](#) | [Book a 30-minute discussion](#)