The standards you're measured on

Eleven frameworks in one subscription, plus four specialist regimes on request.

Every assessment runs the same way: structured questions, evidence filed against the control it supports, an AI analysis grounded in your own answers, and a certifier-ready evidence pack at the end. What changes between them is the yardstick. **One score per framework, one radar across all of them** — so a board sees a single posture, not ten disconnected reports.

My Hub — Framework Coverage. Every completed assessment on one radar against the 70% threshold, SOC 2 derived from what you have already done. Illustrative data.

INCLUDED — ASSESSMENT BUNDLE

One subscription, eleven frameworks

FRAMEWORK	SCOPE	WHAT IT ANSWERS
Cyber Essentials	41 q · 6 controls	Are we certifiable against the UK baseline, and what blocks it?
NCSC CAF	99 q · 14 principles	Essential-function assurance for OES, GovAssure and NHS DSPT.
ISO 27001:2022	76 q · 8 sections	Ready for Stage 1, Stage 2 or surveillance?
ISO/IEC 42001	68 q · 9 sections	Is our AI management system ready for its next certification audit?
DORA	55 q · 5 pillars	How many of DORA's 45 requirements you meet, and what to fix first.
NIS2 Directive	60 q · 14 themes	How many of the Directive's 43 requirements you meet, and what to fix first.
NIST CSF 2.0	70 q · 6 functions	The Implementation Tier each function reaches, and what lifts it to your target.
CIS Controls v8.1	153 Safeguards	Technical control coverage, scoped to IG1, IG2 or IG3.
SOC 2 Readiness	74 q · AICPA TSC	Where we would stand before engaging an auditor.
UK GDPR	65 q · 12 sections	Legal gates first, maturity second — never one blended score.
Compliance Readiness	30 q · free, no account	Which framework are we closest to? Start here.

SPECIALIST — ACCESS GRANTED PER ENGAGEMENT

Regulated and federal regimes

CMMC Level 2	110 requirements	Your SPRS score, the status your answers qualify for, and what a POA&M can cover.
FedRAMP Rev5 Moderate	323 controls · 18 families	Readiness ahead of a 3PAO assessment, with unmet controls framed as POA&M items.
HIPAA Security Rule	47 q · 45 CFR Part 164	Required vs Addressable scored differently, because the rule treats them differently.
DCC Readiness	149 controls · Levels 0–3	MOD Defence Cyber Certification: the level you meet, and what blocks the next.

Post-quantum readiness is already in four of these. CAF, DORA, NIS2 and the Cyber Resilience Maturity Assessment each ask about cryptographic inventory, supplier migration plans and long-lived data, because harvest-now-decrypt-later is a problem today.

HUB-AUTHORED

Four assessments you won't find anywhere else

Written to answer the questions a published standard never asks.

The eleven frameworks on the previous page are transcriptions: someone else set the questions, and the Hub's job is to score them honestly and evidence them properly. These four are different. Each was authored by Cyber3D to answer a question a buyer, a board or an investment committee actually asks — and none of them has a published standard behind it that we could simply have copied.

Cyber Resilience Maturity

14 THEMES · 73 QUESTIONS · IASME-ALIGNED

"Where do we actually stand — and which certification should we go for next?"

A broad maturity baseline that then derives your readiness for Cyber Essentials, CAF, DORA, ISO 27001, AI Governance and SOC 2, so the next certification is a decision rather than a guess.

SOC Maturity & AI Readiness

7 DOMAINS · 70 QUESTIONS

"Is our Security Operations Centre actually working — and is it ready for AI?"

Benchmarks a SOC across governance, people, process, technology and service delivery on a six-level maturity ladder, with a dedicated AI-readiness domain and a target operating model.

PE Cyber Due Diligence

14 THEMES · 70 QUESTIONS · IASME TIERS

"What cyber risk are we buying, and what will it cost to fix?"

A deal-side assessment producing an IASME tier, a red/amber/green deal-risk rating, a 100-day remediation plan and an estimated cost band — sized for an investment committee, not an auditor.

AI Governance

7 DOMAINS · 56 QUESTIONS

"Can we prove the AI we build and buy is governed?"

What the EU AI Act and data protection law ask of your AI, told apart from ISO/IEC 42001 and NIST AI RMF good practice, plus a domain on the AI you actually run — discovery, supply chain and security posture.

WHAT THEY SHARE

The same machinery underneath

Evidence, not assertions
Every question takes a supporting note and file attachments, held against that question — so the answer is defensible a year later.

Analysis you can pace
Self-Score is free and instant. Standard is one AI pass. Deep adds root-cause per theme, an executive summary and cross-checks against your evidence.

Signals from your own tooling
Optional. Connected Microsoft 365, Okta, CrowdStrike or Google Workspace tenants surface live MFA, patching and detection figures beside the questions they bear on.

A tracked plan, not a PDF
Findings become owned remediation actions with effort estimates, and a risk register entry where the finding is a standing risk.

Cross-framework reuse
Answers already given elsewhere surface as informational cross-references — never auto-filled, always attributed to the assessment they came from.

Re-runnable
Each is a living assessment with reassessment reminders, so movement is tracked continuously rather than rediscovered once a year.

Built to your own framework. Where none of the above fits, we author a bespoke framework to your requirements — your control set, your wording, your scoring — and it runs on the same platform as everything here: evidence, depth tiers, remediation, exports. A commissioned framework is licensed exclusively to the organisation it was built for, and is visible to nobody else.

Multi-scope where it matters. A group with several operating companies, or an MSP running a client portfolio, can keep separate named runs of the same assessment side by side — each with its own evidence, score history and remediation plan.

securitycompliancehub.io

DEEP DIVE 1 OF 4

AI Governance Assessment

7 domains · 56 questions · EU AI Act, UK GDPR, ISO/IEC 42001, NIST AI RMF

Most AI governance work stops at policy. This assessment tests whether the policy is operating — and then asks the question policy never reaches: **do you actually know what AI you are running?** Six questions about your AI come first: EU reach, building or using AI, high-risk use, transparency, general-purpose models and personal data. They decide which questions apply, from 28 to all 56, and which of them the law requires.

Seven domains. D7 is the newest.

THE SEVEN DOMAINS

- **D1 · AI Strategy & Governance** — accountability, board oversight, policy
- **D2 · AI Risk Management** — risk classification and review gates
- **D3 · Data Governance** — training and inference data provenance
- **D4 · Transparency & Fairness** — disclosure, bias testing, human oversight
- **D5 · AI Security & Robustness** — adversarial input, model abuse, resilience
- **D6 · Lifecycle & Accountability** — change control, monitoring, retirement
- **D7 · AI Asset Inventory, Supply Chain & Security Posture**

> **The report counts what the law asks of you** — requirements met, what to fix first, and which duties apply now and which from 2 December 2027 — with the good practice beside it.

DOMAIN 7 — THE AI YOU ACTUALLY RUN

Inventory, supply chain and security posture

A register built from self-declaration misses most real AI use. D7 assesses whether AI is *discovered* — SSO and OAuth app grants, cloud billing and resource tags, network and CASB logs, dependency scanning — and whether what is found is then classified, documented and monitored like any other asset class.

Automatic discovery incl. AI inside SaaS · **Shadow AI** detection and a route back into governance · **Classification** by EU AI Act risk category · **AI bill of materials** for models you train · **Supplier model documentation** and model cards · **Verified weights** — hash or signature, scanned · **Continuous posture** — exposed endpoints, keys, permissions

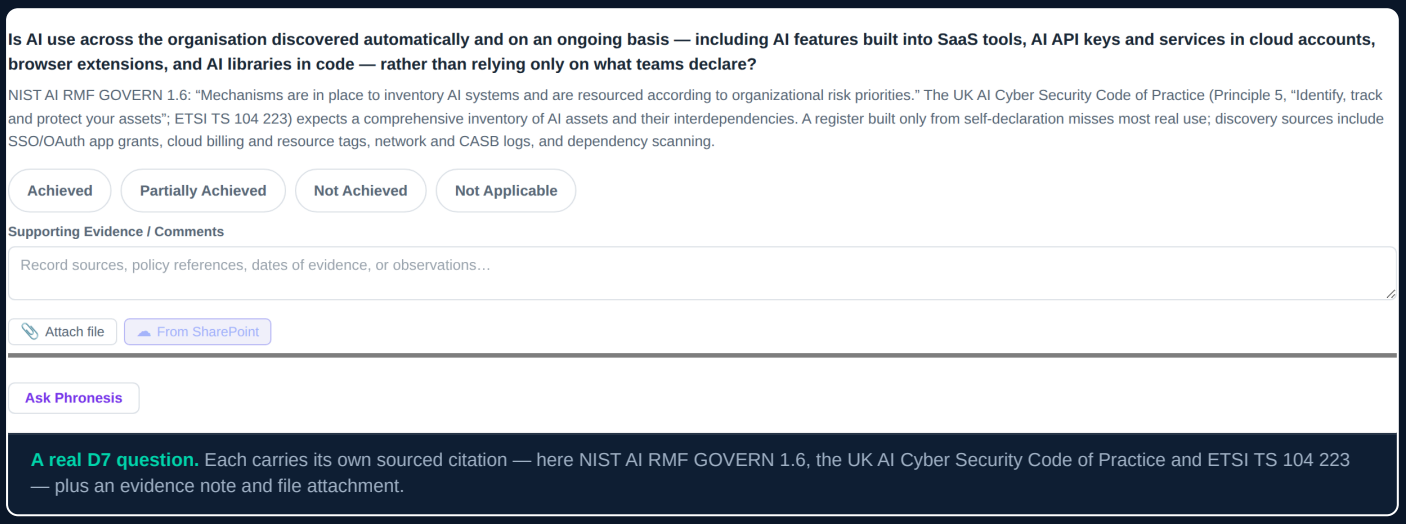

A real D7 question. Each carries its own sourced citation — here NIST AI RMF GOVERN 1.6, the UK AI Cyber Security Code of Practice and ETSI TS 104 223 — plus an evidence note and file attachment.

SOC Maturity & AI Readiness

7 domains · 70 questions · six-level maturity ladder

Organisations invest heavily in Security Operations and still struggle to answer one question: **is it working?** This assessment benchmarks a SOC — in-house, outsourced, hybrid or co-managed — against an industry-wide maturity model, and produces a target operating model rather than a list of tools to buy.

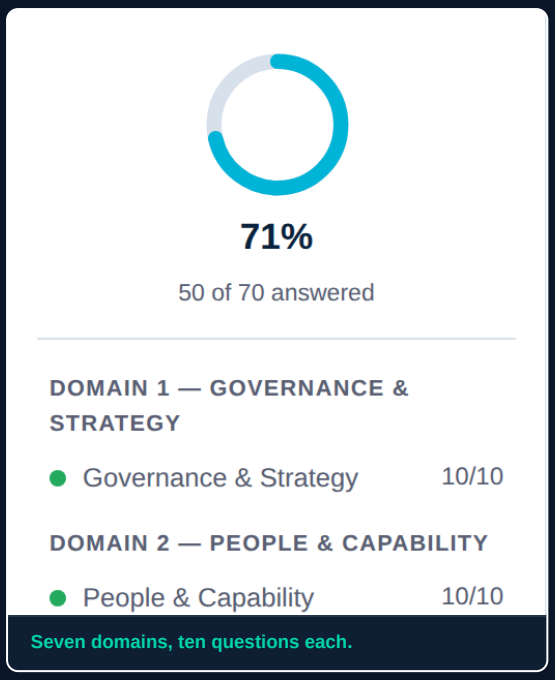

WHAT GETS ASSESSED

- **D1 · Governance & Strategy** — mandate, funding, reporting line
- **D2 · People & Capability** — cover, skills, retention, escalation
- **D3 · Process & Playbooks** — triage, escalation, documented response
- **D4 · Technology & Integration** — telemetry coverage and tuning
- **D5 · Detection & Response Services** — what the SOC actually delivers
- **D6 · AI SOC Readiness** — where automation and AI genuinely help
- **D7 · Target Operating Model** — the shape it should be in 12 months

HOW IT SCORES

A ladder, not a pass mark

A SOC is not compliant or non-compliant — it is somewhere on a path. Each domain and the overall result land on a six-level ladder, so "Level 2 — Developing" reads as a position with a next step rather than a failure.

L0 Non-existent · **L1** Initial / Ad-hoc · **L2** Developing · **L3** Defined · **L4** Managed · **L5** Optimising

Sized to the SOC you have
Enterprise, small/mid-size or MSSP — the question set adjusts, so a 6-person team is not judged against a 24/7 follow-the-sun operation.

Practitioner evidence carries weight
Notes recorded against each question are synthesised into a supporting analysis, so planned work and budget constraints count as partial mitigation.

Built for procurement too
The same framework shapes a greenfield SOC build or an RFP, so requirements and expectations are stated in the same language a bidder will answer in.

Best for leadership teams who need a strategic view of protective capability before committing further spend; organisations tendering for SOC services or weighing in-house against outsourced; and anyone who has to show a regulator, customer or insurer that detection and response actually function.

ON THE SAME PLAN

The full SOC-CMM model

When a SOC wants the industry's reference model rather than this shorter assessment, the same plan opens the complete **SOC-CMM 2.4.2** (advanced): Rob van Os's SOC Capability Maturity Model, used under CC BY-SA 4.0. **667 questions** across five domains (Business, People, Process, Technology and Services) and 27 aspects. Each question is weighted by how much it matters to your SOC, and tooling or services you don't run are left out. Maturity (0 to 5) and capability (0 to 3) are worked out exactly as SOC-CMM's own workbook works them out; the report adds a target level, the fewest questions to raise to reach it, and a roadmap.

PE Cyber Due Diligence

14 themes · 70 questions · IASME tier and deal-risk rating

Cyber diligence on a transaction has a different job from a compliance assessment. It has to price risk, not certify control — and it has to do it on the timetable of a deal. This assessment produces **a tier, a deal-risk colour, a 100-day plan and a cost band**, in language an investment committee can act on.

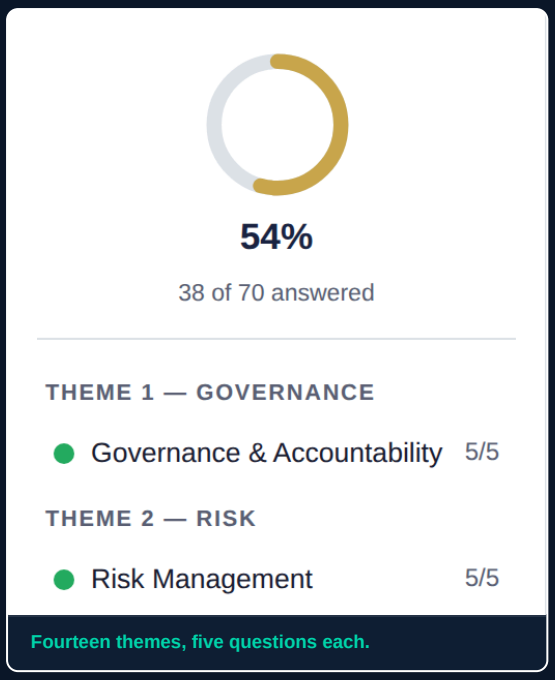

WHAT THE OUTPUT GIVES AN IC

- **An IASME tier** — Tier 1 Initial through Tier 4 Managed, from the weighted theme scores
- **A deal-risk RAG** — red at Tier 1, amber at Tier 2, green at Tier 3 and above
- **A 14-axis profile** — so a single weak theme is visible, not averaged away
- **A 100-day remediation plan** — what the first post-completion quarter has to cover
- **An estimated cost band** — enough to argue a price chip or a warranty

> **Scaled to the target.** A micro or small target is assessed on 46 of the 70 questions — the 24 enterprise-only items (SIEM, PAM, dedicated SOC, board reporting cadence) are scored not-applicable rather than counted as failures.

DESIGN DECISIONS THAT MATTER ON A DEAL

Built for the diligence room

Breach disclosure is reverse-scored
"Has there been a disclosed breach?" is asked as a plain yes/no, and a yes scores zero — the one question where the usual maturity vocabulary reads ambiguously in both directions.

Evidence travels with the finding
Data-room documents attach to the theme they evidence, so the diligence file stands up when it is handed to the buyer's insurer or lender.

Same themes as the maturity baseline
T1–T14 match the Cyber Resilience Maturity Assessment, so a portfolio company can be re-baselined post-completion against the same structure it was bought on.

> **Best for** private equity and corporate development teams running pre-completion diligence; portfolio operations teams re-baselining an acquired company; and vendors preparing for sale who would rather find the issues before a buyer does.

securitycompliancehub.io

DEEP DIVE 4 OF 4

Cyber Resilience Maturity Assessment

14 themes · 73 questions · the starting point for everything else

Most organisations arrive knowing they need "a framework" without knowing which one. CRMA answers that first: a broad maturity baseline across 14 themes, which then **derives your readiness for six named frameworks** so the next certification is chosen on evidence rather than on whichever standard came up in a meeting.

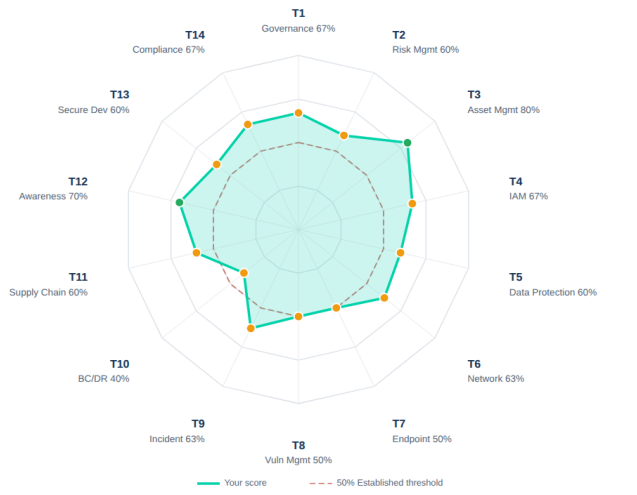

14-axis maturity profile against the 50% Established threshold. Illustrative data.

FOUR MATURITY TIERS

Advanced 75%+ · **Established** 50–74% · **Developing** 25–49% · **Foundation** under 25%

WHAT YOU GET BACK

- **Framework readiness** for CE, CAF, DORA, ISO 27001, AI Governance and SOC 2, with one recommended next step
- **A 90-day roadmap** — four milestones across twelve weeks, not a backlog
- **Critical gaps and strengths** separated, so the report is usable at board level
- **Risk scenarios** mapping your weak themes to real-world loss events

Scaled to the organisation. Micro and small organisations answer 48 of the 73; the 25 enterprise-only questions are scored not-applicable.

POST-QUANTUM READINESS

Assessed now, because the exposure starts now

Data captured today can be decrypted once a cryptographically relevant quantum computer exists, so anything with a long confidentiality life is already exposed. CRMA's Data Protection theme asks three dedicated questions — **where public-key cryptography is used**, **whether suppliers have a migration plan**, and **whether long-lived data is prioritised first**. CAF, DORA and NIS2 carry equivalents, and the Hub watches NIST and NCSC post-quantum publications weekly.

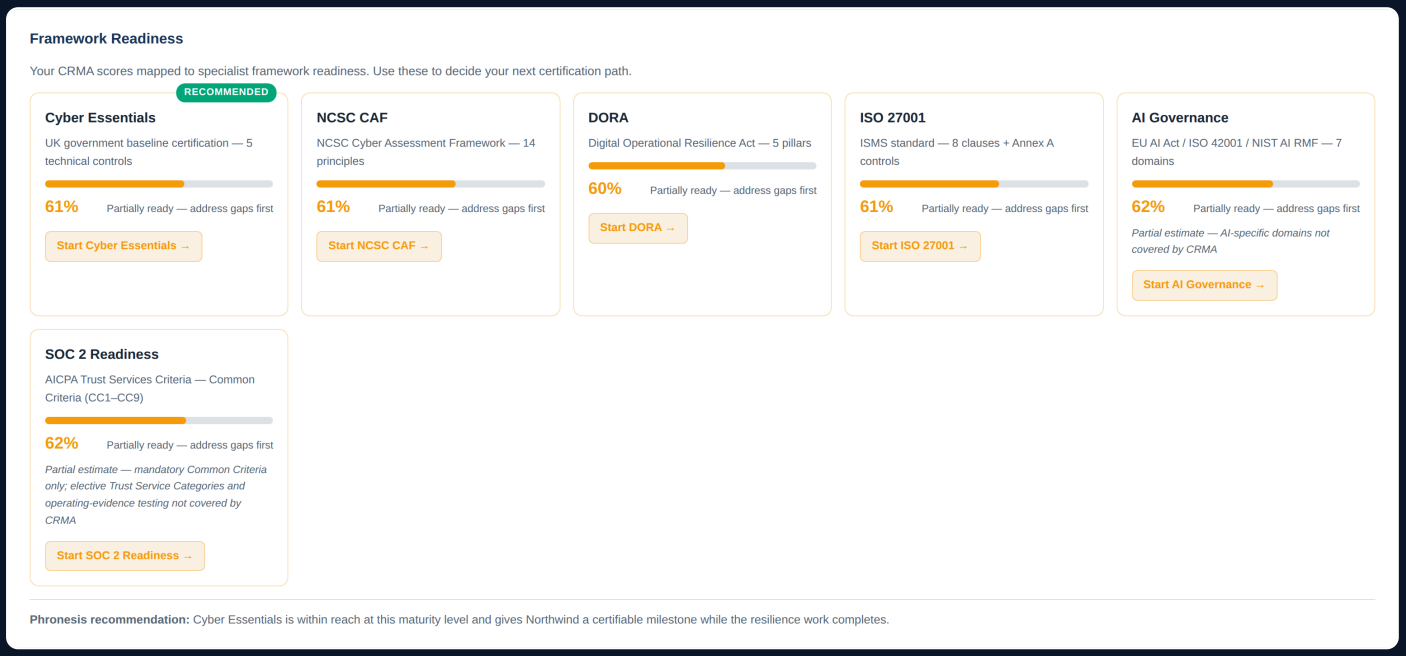

Screenshots show the real product interface populated with illustrative data for one fictional organisation; no customer data is shown. Assessments are readiness and maturity reviews — they are not certifications, audits or legal determinations, and they do not replace assessment by a certification body, a 3PAO, an auditor or a Competent Authority. Question counts are current at the date of issue. Cyber3D Ltd.