

SCORE IT QUARTERLY • 25 QUESTIONS • 20 MINUTES

AI security self-assessment.

Twenty-five questions across five domains. Score each honestly: **0** = no • **1** = partly / informally • **2** = yes, and I could show you. "Could show you" is the test — if the evidence is a feeling, score 1.

A • VISIBILITY — DO YOU KNOW WHAT YOU RUN?		
—	A1. We have a written list of AI tools in use, updated in the last quarter Including the unofficial ones staff declared	0 • 1 • 2
—	A2. We know which tools touch customer or staff personal data Flagged on the list, not held in someone's head	0 • 1 • 2
—	A3. Staff have a safe route to declare new tools — and use it At least one declaration in the last quarter is the evidence	0 • 1 • 2
—	A4. We check expenses and the admin console for new AI apps quarterly Diarised, done, dated	0 • 1 • 2
—	A5. We know where AI features are switched on inside software we already own CRM, office suite, helpdesk — the AI you didn't choose	0 • 1 • 2
B • POLICY & PEOPLE — DO HUMANS KNOW THE RULES?		
—	B1. A written AI usage policy exists, two pages or fewer, under a year old Length is part of the score — forty pages scores 1	0 • 1 • 2
—	B2. Every member of staff has seen it and could state the data rule Ask two people the three bins; that's the audit	0 • 1 • 2
—	B3. There is an approved-tools list that says yes to something A policy that only prohibits has already been routed around	0 • 1 • 2
—	B4. One named person approves new tools, within a stated time A name and a week, not "IT" and "eventually"	0 • 1 • 2
—	B5. Near-misses get reported — and thanked, not punished At least one reported near-miss a year is healthy, not alarming	0 • 1 • 2

C • DATA PROTECTION — IS THE VALUABLE STUFF STAYING IN?		
—	C1. Work AI use is on company accounts, not personal ones Business tiers, centrally managed	0 • 1 • 2
—	C2. Training-on-our-data is switched off, and we kept evidence The screenshot with a date on it	0 • 1 • 2
—	C3. AI accounts are tied to joiners/leavers like email is SSO or an offboarding checklist line — leavers lose access same day	0 • 1 • 2
—	C4. RED-bin data has never knowingly gone into an AI tool — or was handled as an incident when it did Both halves count; the second half is the mature answer	0 • 1 • 2
—	C5. Meeting recording/transcription by AI requires telling the room And external parties are asked, not surprised	0 • 1 • 2
D • FRAUD & SUPPLIERS — THE OUTSIDE WORLD		
—	D1. Payment changes are verified by callback to a number already on file Written procedure, no exceptions clause for seniority	0 • 1 • 2
—	D2. The team has been told, this year, that polish no longer proves legitimacy Training that mentions deepfakes and AI phishing specifically	0 • 1 • 2
—	D3. We have run one fraud tabletop in the last year 45 minutes, three people, the deepfake-payment scenario	0 • 1 • 2
—	D4. Key suppliers have answered our four AI questions in writing What touches our data • training • off-switch • breach notification	0 • 1 • 2
—	D5. Customer contracts and questionnaires mentioning AI get a considered answer Someone owns the answer and it matches reality	0 • 1 • 2
E • AGENTS & AUTOMATION — SOFTWARE THAT ACTS		
—	E1. We know whether any AI here can act — send, change, delete, pay "None" scored with confidence is a 2; "not sure" is a 0	0 • 1 • 2
—	E2. Every agent has a named owner and its own identity Its own account and credentials — never a human's login	0 • 1 • 2
—	E3. Consequential actions require human approval Money, deletions, bulk and external sends gated on a person	0 • 1 • 2
—	E4. We can revoke an agent's access in under a minute, and have tested it The kill switch drill, run once	0 • 1 • 2
—	E5. Connectors (MCP/plugins) come from trusted sources with minimal scopes, reviewed quarterly Unused connectors removed; permissions read before granting	0 • 1 • 2

SCORE RECORD

Track the trend, not the number

Score quarterly. A flat good score beats a spiky brilliant one; what you are looking for is zeros becoming ones and ones becoming twos, one quarter at a time.

SCORING

TOTAL / 50	WHERE YOU ARE	DO NEXT
0–15	Exposed. AI is in use; governance isn't. You are not unusual, but you are relying on luck.	Start the roadmap (01) at phase 1 — this week. First win is the amnesty inventory.
16–30	Foundations forming. Policy or inventory exists; the gaps cluster in whichever domain scored lowest.	Take your weakest domain and close its zeros first. Zeros are cheaper to fix than ones.
31–42	Governed. Better than most firms of any size. The remaining risk is drift.	Diarise the quarterly hour; put the score in the board briefing (05) each quarter.
43–50	Exemplary. Verify with fresh evidence — a score this high is occasionally a scoring problem.	Re-score with a sceptical colleague marking. Trend, not number.

QUARTER	A /10	B /10	C /10	D /10	E /10	TOTAL /50	ONE ACTION AGREED
[Q]							
[Q]							
[Q]							
[Q]							
[Q]							
[Q]							

Rules of engagement: score with evidence open, not from memory.
"Could show you" means within five minutes, to an outsider. One agreed action per quarter beats five aspirations. If a score goes DOWN, write down why – that line is the most useful one this document will hold.

© Peter Bassill · peterbassill.com · CC BY 4.0. Companion documents: 90-day roadmap (01) · staff cheat sheet (02) · usage policy template (03) · agents & MCP guide (04) · board briefing template (05) · DPIA guide (07). The full kit lives at peterbassill.com/ai-security-starter-kit.