Principle Security.

MATURITY PLAYBOOK

Your maturity score is a snapshot. Here's how to *move it*.

The companion playbook to the Principle Security Maturity Scorecard — what each NIST CSF 2.0 function looks like when it's working, the quick wins that move scores fastest, and how to turn a self-assessment into a board-ready roadmap.

2026 EDITION · PRINCIPLESEC.COM

CONTENTS

What's inside

ORIENTATION

How to read your score

The scorecard rates fifteen practices across the six NIST CSF 2.0 functions on a four-point operating scale — from *not in place* to *managed & measured*. Your overall percentage maps to a tier:

Tier	What it means in practice
Tier 1 · Initial (0–29%)	Ad-hoc and reactive. The good news: the highest-impact fixes from here are also the cheapest — fundamentals, not platforms.
Tier 2 · Developing (30–54%)	Core controls exist; coverage is uneven and mostly undocumented. The work is consistency and evidence, not new tools.
Tier 3 · Defined (55–79%)	A real, documented program. The gap to Tier 4 is measurement: testing what you run and proving what you claim.
Tier 4 · Managed (80–100%)	Operating and monitored. Focus shifts to optimization, quantification (see our FAIR guide), and demonstrating trend to leadership.

Three rules before you act on it

- **The profile matters more than the average.** A 60% overall with Recover at 20% is a specific, urgent story; the average hides it. Work your two lowest functions first.
- **Score what you can evidence.** If nobody can produce the access-review record, the honest answer was "partial" — re-score with that discipline and the roadmap writes itself.
- **Maturity is cadence, not architecture.** Almost every question in the scorecard is really asking: does this happen on a schedule, and can you prove it happened?

FUNCTION 1 OF 6

GV · Govern — ownership and oversight

What it's really asking: does someone with authority own security, does leadership see cyber risk on a cadence, and are the rules written down and current?

What good looks like

- ✓ A named accountable owner — CISO, vCISO, or a leader with explicit charter and budget authority
- ✓ Cyber risk on the board/leadership agenda quarterly, with minutes to prove it
- ✓ Policies approved, dated within 12 months, and short enough that people have actually read them
- ✓ A written risk appetite the organization can test decisions against

Fastest score movers

- **Name the owner in writing** — a one-page charter signed by the CEO changes this function's score in an afternoon.
- **Put a standing item on the board agenda** — the cadence, once calendared, tends to survive.
- **Date-stamp the policy review** — reviewing and re-approving existing policies counts; perfection doesn't, currency does.

> **Evidence to keep:** charter, board minutes mentioning cyber, policy approval log, risk appetite statement. Four documents — that's the whole GV audit trail.

FUNCTION 2 OF 6

ID · Identify — knowing what you protect

What it's really asking: can you produce a current list of your hardware, software, data, and vendors — and do you know which of them can hurt you?

What good looks like

- ✓ An asset inventory that updates itself (agent- or network-discovered), not a spreadsheet from the last audit
- ✓ Sensitive data locations mapped — where member/customer/regulated data actually lives, including the SaaS nobody registered
- ✓ Vendors tiered by data access and criticality, with the top tier reviewed on a schedule

Fastest score movers

- **Turn on discovery you already own** — RMM, EDR, and M365/Google admin consoles can produce 90% of the inventory today.
- **Tier the vendor list in one sitting** — three buckets by data access. The tiering *is* the program; depth comes later.
- **Ask finance for the SaaS spend report** — the fastest shadow-IT discovery tool ever invented.

> **Why this function multiplies the others:** every Protect, Detect, and Recover score is capped by what you know exists. Unknown assets are unprotected, unmonitored, and unrecoverable — by definition.

FUNCTION 3 OF 6

PR · Protect — the fundamentals that prevent

What it's really asking: MFA everywhere it matters, least-privilege access with reviews, patched endpoints on a schedule, and people trained against the attacks that actually arrive.

What good looks like

- ✓ MFA enforced on email, remote access, and *every* privileged account — with documented, expiring exceptions
- ✓ Quarterly access reviews that produce a record (and occasionally actually remove someone)
- ✓ Endpoints centrally managed; patch SLAs by severity, measured monthly
- ✓ Awareness training on a calendar, with phishing simulations that report click *and* report rates

Fastest score movers

- **Close the MFA gap list** — the audit usually finds 3–5 privileged paths (core admin, hypervisor, backup console, vendor portals) still on passwords. Highest risk-per-hour fix in this playbook.
- **Run one access review now** — a dated export with sign-off converts "we do reviews informally" into evidence.
- **Report the patch number monthly** — measurement alone tightens the cadence.

> **Watch for:** Protect is where scores are most commonly over-claimed. "MFA is enabled" and "MFA coverage is verified across all privileged access paths" are different maturity levels — the difference is a coverage report.

FUNCTION 4 OF 6

DE · Detect — finding trouble quickly

What it's really asking: if an attacker were inside right now, would anything tell you — and would anyone be listening when it did?

What good looks like

- ✓ Security logs centrally collected from endpoints, identity, email, and network — with retention that survives an investigation
- ✓ Alerting tuned to behavior (impossible travel, privilege escalation, mass file changes), not just signatures
- ✓ A defined human answer to "who sees the alert at 2 a.m., and what do they do?" — in-house or MDR

Fastest score movers

- **Centralize identity and email logs first** — most intrusions touch both early; M365/Google logging upgrades are configuration, not procurement.
- **Tune the top five alerts** — five high-fidelity behavioral alerts someone actually triages beat five hundred ignored ones. Alert fatigue is a maturity *reducer*.
- **Decide the 2 a.m. question honestly** — if the answer is "nobody," an MDR service is the fastest legitimate path from Tier 1 to Tier 3 in this function.

> **The metric that matters:** time from first malicious action to first human awareness. Industry dwell times are measured in days-to-weeks; every hour you remove compresses the entire incident that follows.

FUNCTION 5 OF 6

RS · Respond — acting under pressure

What it's really asking: when detection fires, do people know their roles, and have they practiced — or will the first rehearsal be the real thing?

What good looks like

- ✓ A documented IR plan with named roles, escalation criteria, and out-of-band contact paths
- ✓ A tabletop exercise in the last 12 months with written after-action items — and evidence those items got done
- ✓ Breach materiality and notification obligations mapped *before* the incident (regulators, insurers, customers, law enforcement)
- ✓ Pre-established relationships: IR retainer or firm-on-call, breach counsel, insurer's hotline in the plan

Fastest score movers

- **Run a 90-minute tabletop this quarter** — ransomware scenario, leadership in the room. One exercise moves both RS statements and produces your most persuasive evidence artifact.
- **Write the notification decision tree** — who decides "material," on what clock, notifying whom. One page; legal reviews it once.
- **Refresh the contact card** — printed, wallet/fridge-grade, because the incident that matters may take your email with it.

FUNCTION 6 OF 6

RC · Recover — getting back to business

What it's really asking: are your backups the kind ransomware can't reach, and have you actually restored from them — timed, recently, on purpose?

What good looks like

- ✓ Backups automated, **isolated or immutable** (an attacker with admin credentials cannot delete or encrypt them), and covering SaaS data — not just servers
- ✓ Restores tested on a rotation: timed, documented, with failures treated as findings and fixed
- ✓ RTO/RPO defined per critical system *with the business*, and the tested numbers compared against them
- ✓ Recovery communication planned — who tells customers/members what, when systems are down

Fastest score movers

- **Time one restore this month.** The single most valuable maturity action in this entire playbook. It either proves your resilience or finds the problem while it's still cheap.
- **Flip on immutability** — most modern backup platforms have object-lock/immutability as a setting, not a purchase.
- **Write RTO/RPO for your top five systems** — one workshop with operations; the numbers don't need to be flattering, they need to exist.

> **Why RC scores lowest almost everywhere:** recovery only gets exercised by disasters unless someone schedules the rehearsal. Scheduled rehearsal is the entire difference between Tier 1 and Tier 3 here.

THE ROADMAP

From score to roadmap: 90 / 180 / 365

Sequence by risk-per-effort, not by function order. The pattern that works for nearly every Tier 1–2 organization:

90 days — stop the bleeding, start the evidence

- ✓ MFA gap closure on privileged paths (PR) · one timed restore + immutability on (RC)
- ✓ Named owner + board cadence started (GV) · one access review with a record (PR)
- ✓ 90-minute tabletop (RS) — doubles as leadership buy-in for everything below

180 days — build the operating rhythm

- ✓ Automated asset discovery live; vendor list tiered (ID) · identity + email logs centralized, top alerts tuned or MDR selected (DE)
- ✓ Policy suite reviewed and re-approved (GV) · patch SLA measured monthly (PR)
- ✓ IR plan updated from tabletop findings; notification tree written (RS)

365 days — measure and prove

- ✓ Quarterly cadence running on rails: access reviews, restore rotations, board packets
- ✓ Re-take the scorecard — expect two functions up a full tier if the cadence held
- ✓ Graduate the conversation: quantify your top scenarios in dollars (see our FAIR Board Risk Guide) and validate controls with a penetration test

REPORTING

Reporting maturity to the board

A maturity score is the rare security metric a board understands instantly — six bars and a trend arrow. Use that, and avoid the traps:

Do

- ✓ **Show the six functions, not the average** — the shape tells the story, and the weakest bar justifies the budget line
- ✓ **Trend it quarterly on the same scale** — improvement a director can see beats any narrative
- ✓ **Tie each planned investment to the bar it moves** — "this $30K moves Recover from Tier 1 to Tier 3" is an approvable sentence
- ✓ **Pair the score with evidence highlights** — restore time, tabletop date, MFA coverage % — so maturity reads as operational fact, not self-grade

Don't

- **Don't chase 100%.** Tier 4 everywhere is rarely the right spend — say where you intend to sit and why (that statement *is* risk appetite).
- **Don't re-baseline quietly.** If scoring gets more honest and numbers drop, say so — credibility outlasts any single quarter's optics.
- **Don't present maturity as risk.** Maturity says how well the machine runs; risk says what it costs when it fails. Boards eventually need both — that's the FAIR conversation.

HOW PRINCIPLE SECURITY HELPS

Want the validated version, with a *roadmap attached*?

The scorecard is directional. Our maturity assessment validates it against evidence, maps results to your regulatory obligations (FFIEC, HIPAA, SOC 2, CMMC), and delivers the sequenced 90/180/365 roadmap with budget-ready justification.

Prefer ongoing leadership? Our vCISO service runs the whole cadence — reviews, tests, tabletops, and board reporting — as a fraction of a full-time hire.

Book a scoping call → principlesec.com/contact

info@principlesec.com · +1 (877) 886-0677
principlesec.com/services/vciso