SAUTERA · FIELD GUIDE

The SPRS Self-Assessment Evidence Checklist

DETECT
↓
DECIDE
↓
ACT
↓
PROVE
↓
IMPROVE

What you must be able to produce - not just attest

Where infrastructure trust is determined.

The SPRS Self-Assessment Evidence Checklist

What you must be able to produce — not just attest

The C3PAO deadline moved. Your self-assessment didn't.

By Joe Augustine · SAUTERA · August 2026

In July 2026, the Department of Defense suspended the CMMC Phase 2 third-party (C3PAO) certification deadline that had been set for November 10, 2026 — along with Phases 3 and 4. A lot of contractors read that as a reprieve.

It isn't. What remains in force is the part with your name on it:

- **Phase 1 still binds.** New DoD solicitations still condition eligibility on a current CMMC Level 1 or Level 2 **self-assessment** and a score posted to the **Supplier Performance Risk System (SPRS)**.
- **Primes are not waiting.** Flow-down demands from prime contractors continue regardless of the federal timeline — for many suppliers, the prime's deadline is the real one.
- **Your SPRS score is a federal representation.** You post it under the shadow of the False Claims Act. The legal risk attaches to the number you post today, not to the certificate you don't yet need. A generous score with no evidence behind it is the worst position in this industry — worse than an honest negative score with a remediation plan and a date.

This checklist covers the question the suspension didn't change: **for each of the 110 NIST SP 800-171 practices you are scoring yourself on, what evidence must you actually be able to produce?**

> **This is educational material, not legal advice.** *Assessment methodology and CMMC program rules change; verify current requirements against the DFARS clauses in your contracts (252.204-7019/-7020/-7021) and with counsel before posting or revising a score.*

How to use this checklist

For every practice family below, ask three questions:

1. **Do we do this?** (The attestation — where most self-assessments stop.)
2. **Can we produce evidence that it operated over time?** (What an assessor — self, prime, or C3PAO — actually accepts: records, not recollections.)
3. **Who owns the artifact, and how fresh is it?** (Evidence with no owner goes stale silently.)

A practice you perform but cannot evidence scores the same, until someone asks. Then it doesn't.

The 14 families — what evidence looks like

`Where continuous, device-layer evidence carries the weight`

These six families live substantially at the device layer. They are also where evidence rots fastest — a screenshot of patch status is stale the day after it's taken. The standard here is **records generated continuously, not assembled**

retrospectively.

SI — System & Information Integrity (3.14.1–3.14.7)

- ☐ Flaw-remediation records per device: what was missing, when it was identified, when it was remediated — not a patch policy PDF.
- ☐ Malicious-code protection presence verified on each endpoint (AV/EDR observed running, not "deployed per policy").
- ☐ Patch currency measured on a defined cadence, with the laggards named. *Evidence standard: a per-device record that shows the state moving over time.*

CM — Configuration Management (3.4.1–3.4.9)

- ☐ Baseline configurations documented — and drift from them detected and dated.
- ☐ Host firewall state and exposed listening surface recorded per device.
- ☐ Change records tying configuration changes to approvals. *Evidence standard: when a baseline slipped, your records show when, on which host, and what happened next.*

SC — System & Communications Protection (3.13.x, esp. 3.13.11, 3.13.16)

- ☐ Disk-encryption posture observed per device holding CUI — not inferred from a deployment policy.
- ☐ Cryptography inventory: where FIPS-validated modules are required, which modules, validated against the CMVP list.
- ☐ Boundary-protection architecture documented and current. *Evidence standard: per-device encryption state with a timestamp, refreshed as the fleet changes.*

RA — Risk Assessment (3.11.1–3.11.3)

- ☐ Vulnerability scans on a defined cadence, with findings tracked to closure — the scan report AND the disposition.
- ☐ End-of-life exposure named: operating systems and software past or approaching end of support are themselves findings. (Windows 10 went end-of-life October 2025 — machines still on it without ESU are accumulating unpatched CVEs monthly, and an assessor will treat them exactly that way.) *Evidence standard: a ranked, dated exposure list — not a raw scanner export nobody triaged.*

AU — Audit & Accountability (3.3.1–3.3.9)

- ☐ Audit logs created, retained per your stated retention, and protected from modification.
- ☐ Evidence of review: who looked, when, what was flagged.
- ☐ Tamper-evident records for security-relevant events — findings, decisions, and fixes. *Evidence standard: the audit trail of the audit trail. Logs that were never reviewed evidence only their own existence.*

CA — Security Assessment (3.12.1–3.12.4)

- ☐ The self-assessment itself, dated, with per-practice scoring rationale.
- ☐ A current System Security Plan (SSP) — the document the entire self-assessment hangs on. No SSP, no score.
- ☐ Plans of Action & Milestones (POA&Ms) with dates for every practice scored as not-met.
- ☐ **Continuous monitoring** records (3.12.3): proof that control effectiveness is watched between assessments, not only during them. *Evidence standard: the assessment is an export from a running system, not an annual project.*

Where the evidence is people, paper, and process

No monitoring platform observes these — including ours. They are self-assessment items in the fullest sense, and they are where honest self-scoring most often finds its negative points.

AC — Access Control: account inventory, least-privilege rationale, CUI-flow documentation, remote-access controls, session records. **IA — Identification & Authentication:** MFA coverage records (who is enrolled, who is exempted and why), password-policy enforcement evidence, credential-lifecycle records. **IR — Incident Response:** the plan, the reporting path (including DFARS 72-hour DIBNet reporting), and — the one everyone misses — evidence you **tested** it (tabletop dates, findings). **AT — Awareness & Training:** completion records per person per year, role-based training for privileged users, insider-threat awareness. **PS — Personnel Security:** screening records, termination checklists showing access revoked, with dates. **PE — Physical Protection:** visitor logs, physical-access records for spaces where CUI lives. **MA — Maintenance:** maintenance records, sanitization of equipment leaving the boundary, supervision records for third-party maintainers. **MP — Media Protection:** CUI media inventory, sanitization certificates, encrypted-transport records for media in transit.

The scoring trap

The DoD Assessment Methodology starts you at **110** and subtracts weighted deductions — **1, 3, or 5 points** per practice not met. The floor is **−203**. Two things follow:

1. **Most first honest scores are negative.** That is normal, expected, and defensible — *if* every deduction carries a POA&M with a date. What is not defensible is a 110 you cannot evidence.
2. **Not everything can ride a POA&M.** Under the CMMC Level 2 rules, the highest-weighted practices generally cannot be deferred to a POA&M at assessment time, and conditional status requires meeting a minimum score threshold. The details here move — verify the current POA&M-eligibility rules against the CMMC program documentation before you rely on them.

The strategic posture: **post the number your evidence supports.** Then improve the evidence, not the number.

Continuous evidence vs. the fire drill

Every path through this checklist ends at the same fork:

- **The fire drill:** each time a score is due — self-assessment, prime demand, or an eventually-returning Phase 2 — someone spends weeks assembling screenshots, exports, and recollections into a folder that is stale on arrival.
- **Continuous monitoring:** the evidence exists because the system that watches your devices writes tamper-evident records as a side effect of operating. The assessment becomes an export. NIST SP 800-137 calls this Information Security Continuous Monitoring; practice 3.12.3 asks for it by name.

The suspension of the C3PAO deadline changed *when* someone else will grade your evidence. It did not change *whether*. Contractors who use the extra time to stand up continuous evidence will export a package whenever the clock restarts. Contractors who treat it as a reprieve will run the fire drill — again, under more scrutiny.

Where SAUTERA fits — and where it doesn't

SAUTERA is an infrastructure-trust platform: it continuously observes each device (Windows and Linux sensors, agentless SSH/WMI/SNMP, and read-only cloud connectors), concludes an honest trust verdict per device — end-of-life

software can never score Trusted here — and writes every finding, decision, and human-approved fix to a tamper-evident record. For the six device-layer families above, that means the evidence exists because monitoring operated, and it exports on demand, including a Continuous Monitoring Attestation for the periods where the news is that nothing went wrong.

What SAUTERA is not: a certification, a substitute for your SSP, or coverage for the people-and-paper families — no software observes your visitor logs or your training records, and a vendor claiming otherwise is selling the fire drill back to you.

See your own device-layer evidence this week: the free tier scores 10 devices — no card, no expiry.
sautera.com/solutions/eol-exposure · questions: sautera.com/contact

---