

TOOLKIT COMPANION

SOC 2 Readiness Guide

The 33 Security criteria in plain English, what evidence usually looks like, and questions for any vendors without their own SOC 2 report.

SCOPE

CC1.1 → CC9.2

PART 1

33 criteria

PART 2

35 vendor questions

How to use the toolkit

You need policies and controls. Policies set the rules, controls put them into practice, and evidence shows the controls operate.

The workbook has six tabs. Start with How to Use, then work the other five in order.

These are criteria, not a prescribed list of controls. Your controls, and the evidence you retain, are based on what is in scope for your environment: your systems, relevant risks, existing processes, and resources.

Start	How to Use	A quick overview of the workbook, the status values, and how the calculations work.
1	Readiness Checklist	Read each criterion, what it looks like in practice, and the evidence to collect. Set Status to In Progress while you gather evidence and Implemented once it is stored. Use N/A only for criteria outside your system boundary.
2	Risk Register	Turn the gaps you care about into rows 11-60. Score Likelihood and Impact from 1 to 5. Critical 20-25, High 15-19, Medium 8-14, Low 1-7. Give every open risk an owner and a date. The example risks below the register are not counted.
3	SOC 2 Report Review	For vendors that will share their report. Record the type, period, opinion (Unmodified, Qualified, Adverse, or Disclaimer), and CPA firm. A trust-center page is where you request the report, not the report itself.
4	Vendor DDQ	For vendors where you do not already have a report to review. Replace the example vendors, send the 35 questions in Part 2, and log each answer as Yes, No, Partial, N/A, or Pending. The answer columns hold one vendor at a time, so duplicate the tab for each additional vendor.
5	Dashboard	Updates as you go. Criteria complete counts rows marked Implemented. Evidence linked counts rows with a note or link. Example rows are never counted.

COMMON GAPS

What a reviewer can actually test

A policy on paper doesn't carry much weight by itself. For a Type 1, auditors want to see that the control is actually in place. For a Type 2, they want proof it operated consistently over the review period. Either way, the evidence is typically artifacts like dated exports, tickets, approvals, and acknowledgments that an independent reviewer can verify.

Areas that are often thin at a first SOC 2:

- ✦ Access reviews that happened but were not recorded
- ✦ Former employees who still have accounts
- ✦ Changes that reach production with no approval trail
- ✦ An incident program nobody has walked through
- ✦ Vendors that were never assessed

These are gaps to close or explain. None of them, by itself, means a qualified or adverse opinion.

PART 1 • READINESS CHECKLIST

The 33 Security criteria

Titles are shortened from AICPA TSP section 100. Each one includes a sample control, what it looks like in practice, and the evidence that shows the control operates. Track status in the Readiness Checklist tab.

CC1 CONTROL ENVIRONMENT

5 CRITERIA

CC1.1 Commitment to integrity and ethical values (COSO Principle 1)

Sample control

Employees acknowledge the code of conduct at hire and every year, and suspected violations go to a named person to investigate.

In practice

People can see what acceptable conduct is. A code of conduct is shared when someone joins and acknowledged again on a cycle you choose. Leadership behavior matches that standard, and deviations are addressed rather than ignored.

Evidence it operates

Code of conduct; onboarding acknowledgments; a re-acknowledgment log; a record that a deviation from the code was addressed.

CC1.2 Board independence and oversight of internal control (COSO Principle 2)

Sample control

Leadership, or a board or advisory group, reviews security status and open risks at least twice a year and records decisions.

In practice

Someone outside day-to-day management actually looks at security and risk. An independent director is one structure. An advisory board is a common early-stage approach. What matters is documented oversight, not the label on the group.

Evidence it operates

Minutes or notes from board or advisory discussions that cover security; a short governance or security charter; materials management shared with that group on risk.

CC1.3 Structures, reporting lines, and authorities (COSO Principle 3)

Sample control

An org chart and a named security owner exist, and security responsibilities are written into the relevant roles.

In practice

It is clear who owns security and who they report to. The security lead has a defined mandate and enough authority to act, not an informal side job with no reporting line.

Evidence it operates

Current org chart; job descriptions for security-relevant roles; a simple RACI or responsibility matrix.

CC1.4 Attract, develop, and retain competent people (COSO Principle 4)

Sample control

New hires get background checks where the role warrants and complete security training within 30 days and then every year.

In practice

People who design, operate, or administer the system are hired and trained for that work. Security awareness training is assigned and tracked. Background checks are a common hiring control where the role warrants them.

Evidence it operates

Training assignment and completion records; hiring checks appropriate to the role (a background-check policy is typical); onboarding steps that include security; role descriptions that name security responsibilities.

CC1.5 Accountability for internal control responsibilities (COSO Principle 5)

Sample control

Security responsibilities appear in role descriptions and performance reviews, and policy violations follow a documented process.

In practice

People are answerable for the controls they own. Control failures are tracked through to a correction, not only written down. There is a known consequence for serious violations of the code of conduct.

Evidence it operates

Policy-violation or disciplinary procedure; remediation tracking for control failures; notes from a periodic control review that show owners and outcomes.

CC2 COMMUNICATION AND INFORMATION

3 CRITERIA

CC2.1 Obtain and use relevant, quality information (COSO Principle 13)

Sample control

An up-to-date system inventory, data-flow diagram, and log sources show what the system is and what it records.

In practice

The information used to run internal control is relevant and reliable enough to act on: timely, accurate, complete enough for the decision, protected, and retained. Typical inputs are an asset inventory, access lists, risk ratings, and incident records that people actually use.

Evidence it operates

Inventory of systems and data used for control decisions; source and owner for the security reports you rely on; an example of a decision that used that information (an access review, a risk rating, or incident triage).

CC2.2 Internal communication of control information (COSO Principle 14)

Sample control

Security policies are published internally and employees acknowledge them, with changes announced to staff.

In practice

Employees can find the security policies and objectives that apply to them, understand their responsibilities, and acknowledge them. Policies are versioned and reviewed on a defined cycle. Annual review is common; the criterion does not set a cadence.

Evidence it operates

Information security policy with version and review date; acknowledgment records; the place employees actually open the policy; a record that a material policy change was communicated.

Communication with external parties (COSO Principle 15)**Sample control**

Customers receive security commitments in contracts, can reach a published security contact, and are notified of relevant changes or incidents.

In practice

Customers, vendors, and other outsiders are told what they need in order to meet their responsibilities, and how to report a security concern. Contract terms usually carry the security obligations. A data processing agreement is typical when personal data is in scope. A public status page is a common way to talk about uptime; it belongs with availability commitments, not with these security criteria.

Evidence it operates

A published security contact or other working channel for outsiders to report a concern; customer contract or security-addendum language; vendor data-handling terms.

CC3 RISK ASSESSMENT

4 CRITERIA

Specify objectives clearly enough to identify risk (COSO Principle 6)**Sample control**

Security objectives tied to customer commitments are written down and reviewed every year.

In practice

Security objectives are written clearly enough that you can tell what would count as a risk to them. For a first SOC 2 this is usually the security commitments for the system: who it serves, what must be protected, and what is inside the boundary.

Evidence it operates

Documented system or security objectives; a scope description (services, data, system boundary); a visible link from the risk assessment back to those objectives.

Identify and analyze risks to objectives (COSO Principle 7)**Sample control**

A risk assessment is performed every year and after significant changes, and the results feed a risk register with owners.

In practice

Risks to those objectives are identified and analyzed on a defined cycle, and the analysis is what you use to decide how to respond. The method is written down, not only a list of risks.

Evidence it operates

A dated risk assessment; a short scoring method (this workbook uses Likelihood × Impact on the Risk Register); named risk owners; an example of a priority or control that changed because of the assessment.

CC3.3 Consider the potential for fraud (COSO Principle 8)

Sample control

The annual risk assessment explicitly considers fraud, such as misuse of admin access or tampering with financial data.

In practice

The risk assessment explicitly considers fraud. Typical scenarios are misuse of access, an insider taking data, and someone bypassing a control so the act is not visible.

Evidence it operates

Fraud scenarios in the risk register; notes on segregation of duties, or the compensating review where duties cannot be split; access-review records for privileged roles.

CC3.4 Identify and assess significant change (COSO Principle 9)

Sample control

Significant changes, such as new products, vendors, or a reorganization, trigger a review of the risk assessment.

In practice

Significant change is assessed for new risk: product, infrastructure, vendors, people, and the technology environment. The look can be lightweight. It has to happen, and it has to be visible afterward.

Evidence it operates

A change or vendor-onboarding step that includes a risk look; examples of risk notes on recent significant changes; a short rule for which changes require that look.

CC4 MONITORING ACTIVITIES

2 CRITERIA

CC4.1 Ongoing and separate evaluations of internal control (COSO Principle 16)

Sample control

Controls are reviewed on a schedule, through internal self-assessment or an outside review, and the results are recorded.

In practice

Someone checks whether controls are present and operating, on a mix of ongoing checks and separate reviews. Capturing a log is not the same thing as evaluating the control.

Evidence it operates

Periodic control self-assessment or review notes; a deficiency log opened from those reviews; evidence the result reached a person who can act.

CC4.2 Evaluate and communicate internal control deficiencies (COSO Principle 17)

Sample control

Deficiencies found in reviews are logged as tickets with owners and due dates and tracked to closure.

In practice

When a deficiency is found, it is evaluated, communicated to the people who must fix it (and to senior management and the board when that is your governance), and tracked until it is remediated.

Evidence it operates

Deficiency or remediation log with owners and dates; tickets showing closure; a record that a material gap was communicated to management.

CC5 CONTROL ACTIVITIES

3 CRITERIA

CC5.1 Select and develop control activities (COSO Principle 10)

Sample control

Each significant risk in the register maps to at least one control, and gaps are accepted or remediated with a decision on file.

In practice

For risks that matter to your objectives, you select control activities that bring risk to an acceptable level. The mix usually includes preventive and detective controls, and duties that should be split are split or paired with a compensating review.

Evidence it operates

A control list or matrix mapping risks to activities; examples of segregation of duties, or the compensating review where a split is not practical.

CC5.2 General control activities over technology (COSO Principle 11)

Sample control

Technology controls for access, change, and operations are defined and assigned to owners.

In practice

Technology general controls support the control activities: identity and access administration, infrastructure configuration, logging, and how technology is acquired and changed. These are built into systems where you can, not only written as policy.

Evidence it operates

Identity-provider settings; a configuration or hardening baseline; logging and alerting ownership; evidence these technology controls are reviewed. Multi-factor authentication is a common choice here; the criterion does not name it.

CC5.3 Deploy controls through policies and procedures (COSO Principle 12)

Sample control

Policies are approved by management, reviewed every year, and put into practice through procedures and tooling.

In practice

Policies say what is expected. Procedures say how people carry that out in the day-to-day work. Someone is accountable for each, and they are reviewed when the system or the risk changes.

Evidence it operates

Policies with owners and review dates; runbooks or SOPs for the material control areas; a recent ticket, checklist, or log showing a procedure was followed.

CC6 LOGICAL AND PHYSICAL ACCESS

8 CRITERIA

CC6.1 Logical access security over protected information assets

Sample control

Production access requires SSO with MFA, access is role-based, and data is encrypted at rest.

In practice

Logical access is implemented over the systems and data you protect: assets are identified, users and services are authenticated, and access rules restrict what they can reach. Encryption at rest, with keys protected in generation, storage, use, and destruction, is a common supplement when the risk assessment calls for it. AES-256 is a typical design choice; the criterion does not name an algorithm.

Evidence it operates

System and data inventory used for access decisions; identity-provider and cloud access configuration; authentication requirements. If you encrypt, where keys live and who can use them.

CC6.2 Register, authorize, and remove user access credentials

Sample control

New accounts are created only from an approved access request.

In practice

People and other users you administer do not receive credentials until access is registered and authorized. Credentials are removed when access is no longer authorized, including when someone leaves. Access is reviewed on a cycle you set so leftover access is caught.

Evidence it operates

Access-request records with approval before provisioning; provisioning steps; completed offboarding samples showing credentials removed; the timing you commit to for removal; a periodic access review and the changes it produced.

Authorize, modify, or remove access by role, least privilege, and segregation of duties**Sample control**

Access is reviewed every quarter, and departing employees lose access within one business day.

In practice

Access to data, software, and functions is granted, changed, and removed based on role and job need. Least privilege and segregation of duties are designed in. A role change updates access instead of leaving the old rights in place.

Evidence it operates

Role definitions; access-change records tied to role changes; a periodic review of roles and rules; segregation notes for sensitive functions.

Restrict physical access to facilities and protected information assets**Sample control**

Physical security of the data center is inherited from the cloud provider, confirmed by reviewing its SOC 2 report.

In practice

Physical access to facilities and to protected assets (offices, data-center space you control, backup media, other sensitive locations) is limited to authorized people, changed when roles change, and reviewed. A company whose production systems sit in a cloud provider's data center usually points to that provider's own physical controls and to the provider's SOC 2 report, plus whatever office controls are inside your own boundary. This is the physical-access criterion. Threats from outside the network boundary are CC6.6.

Evidence it operates

Office badge, key, or visitor process if you have a space, and a review of who has access. For a cloud data center, the provider's SOC 2 report from the provider's assurance portal (not a marketing page), showing physical-access controls. Media-storage access if backup media is in your possession.

CC6.5 Make data and software unreadable before physical assets leave your control

Sample control

Devices and storage are wiped or destroyed when retired, and customer data is deleted on a defined schedule.

In practice

Logical and physical protections are not dropped while data or software can still be read. Before a laptop, disk, or other device is disposed of or leaves your control, the data and software on it are rendered unreadable. This is asset disposal, not employee offboarding (offboarding credentials are CC6.2).

Evidence it operates

Disposal or sanitization procedure; records for wiped or destroyed devices; the path for returned laptops and for retired volumes or snapshots.

CC6.6 Protect against threats from sources outside the system boundary

Sample control

Firewalls or security groups restrict inbound traffic, and remote admin access goes through a VPN or bastion with MFA.

In practice

Logical access measures protect the system from threats originating outside its boundary. Typical measures are restricted inbound paths, extra authentication for remote access, protection of credentials in transit beyond the boundary, and boundary protection that is actually monitored. This is not the physical-access criterion (that is CC6.4).

Evidence it operates

Security-group, firewall, or equivalent boundary rules; remote-access authentication settings; a list of external access points and what may pass through them; evidence boundary alerts are monitored.

CC6.7 Restrict and protect transmission, movement, and removal of information

Sample control

Data in transit is encrypted, and moving data to removable media or unapproved destinations is restricted.

In practice

Information is transmitted, moved, or removed only by authorized users and processes, and it is protected while that happens. Encryption or another secured channel for data leaving the environment is the usual control. TLS 1.2 or later is a typical choice for web traffic, not a version the criterion names. Laptops and removable media that hold information need a protection method too.

Evidence it operates

How public endpoints protect data in transit (configuration or certificate inventory). Rules for sending sensitive files; device protection for laptops (MDM or disk encryption is typical); removable-media rules if you use removable media.

CC6.8 Prevent, detect, and act on unauthorized or malicious software

Sample control

Endpoint protection runs on company devices, and only approved software can be installed on production systems.

In practice

Unauthorized or malicious software is prevented where you can, and detected and acted on when it appears. Installation is limited to authorized people, malware controls are maintained, and software is introduced through a defined change path.

Evidence it operates

Endpoint management or an install restriction; anti-malware or EDR coverage and a sample of an alert that was reviewed; the path used to introduce new software.

CC7.1 Detect configuration changes that add vulnerabilities, and new vulnerabilities**Sample control**

Vulnerability scans run on a schedule and after significant changes, and findings are fixed within severity-based timelines.

In practice

You detect two things: configuration changes that introduce vulnerabilities, and newly discovered vulnerabilities the system is susceptible to. Findings are remediated on a timeline tied to severity. Vulnerability scanning, after significant change and on a periodic cycle, is the usual detection method. A web-application penetration test is one way teams often go deeper on this criterion. It is not required evidence.

Evidence it operates

Vulnerability scan results and the remediation that followed; when scans run, including after significant change; a configuration standard and a way you notice drift. A pentest report, if you have one.

CC7.2 Monitor system components for anomalies that may be security events**Sample control**

Logs from production systems are centralized, and alerts for anomalous activity go to an on-call responder.

In practice

System components are monitored for anomalies that indicate malicious acts, natural disasters, or errors that could stop you meeting your objectives. Anomalies are analyzed to decide whether they are security events. Alerts that are generated and never reviewed do not meet the point of this criterion.

Evidence it operates

Logging and detection configuration; the rules or signals you treat as anomalies; samples of triage (a ticket, on-call note, or review log); a named owner for the detection tools.

CC7.3 Evaluate security events to determine whether they are security incidents

Sample control

Alerts are triaged against defined criteria to decide whether they are security incidents.

In practice

Security events are evaluated to decide whether they could or did cause a failure to meet objectives. That failure is a security incident. If it is, action is taken to prevent or address it. This is the evaluation step. The response program itself is CC7.4.

Evidence it operates

Severity or event-to-incident criteria; examples of events that were reviewed and either escalated or closed; the named people who do that evaluation.

CC7.4 Respond to identified security incidents

Sample control

A documented incident response plan assigns roles, and incidents are tracked from detection through resolution.

In practice

Identified security incidents are handled by a defined response program: understand what happened, contain it, remediate it, and communicate with the people who need to know. Roles are assigned, including outside help if you will use it. The program is evaluated on a cycle you choose. Walking through the program (a tabletop or a real incident) is how you learn whether it works. An untested plan is weaker preparation.

Evidence it operates

Incident-response program covering roles, containment, remediation, and communication; a record from a real incident or a tabletop; notes showing a lesson was fed back into the program.

CC7.5 Recover from identified security incidents

Sample control

After an incident, systems are restored and a post-incident review records root cause and follow-up actions.

In practice

After a security incident, you restore the affected environment, determine root cause, communicate what recovery required, and change preventive or detective controls so a repeat is less likely. Testing the recovery steps is a point of focus. Broader business-disruption planning, including keeping the service going through other kinds of disruption, is CC9.1, not this criterion.

Evidence it operates

Recovery steps (rebuild, patch, configuration change, restore); a root-cause note from an incident or exercise; a record of a recovery test if you have run one; evidence lessons changed the plan. Business-continuity scope beyond incident recovery belongs with CC9.1.

CC8 CHANGE MANAGEMENT

1 CRITERION

CC8.1 Authorize, design, test, approve, and implement changes

Sample control

Code changes require pull request review and approval, pass automated tests, and are deployed through a controlled pipeline.

In practice

Changes to infrastructure, data, software, and procedures are authorized, designed or acquired, configured, documented, tested, approved, and implemented. Emergency changes have a path that still gets reviewed. Peer review before production, including protected branches, is a common way software teams implement this. It is a control choice.

Evidence it operates

Change-management policy; a sample change showing test and approval before production; deployment or ticket history; an emergency-change record if that path has been used.

CC9.1 Mitigate risks arising from potential business disruption**Sample control**

Business disruption risks are assessed, and backups are tested by restoring them on a schedule.

In practice

You identify disruption that would stop you meeting your objectives, then select and develop mitigations: procedures, communications, and alternative processing. The points of focus also say to consider insurance to offset financial impact. Cyber insurance is something to consider and, if you buy it, to keep on file. Backup and restore often support the mitigation you select. Incident recovery mechanics stay on CC7.5.

Evidence it operates

Disruption scenarios and the mitigations you selected; continuity or recovery procedures tied to those scenarios; backup approach and a restore test when backups are part of the mitigation. If you use insurance, the certificate. If you considered it and declined, that consideration noted in the risk record.

CC9.2 Assess and manage risks from vendors and business partners**Sample control**

Vendors are risk-ranked at onboarding and reviewed every year using their SOC 2 report or a security questionnaire.

In practice

Vendors and business partners who can affect the system are assessed before you rely on them and again on a cycle you set. Scope, roles, requirements, and service levels are defined. Use SOC 2 Report Review when you have their report, and Vendor DDQ when you do not.

Evidence it operates

Vendor list with a risk tier; assessment records; contract or security terms; the periodic review; completed report reviews or DDQ decisions for vendors that can affect the system.

35 questions for vendors

For vendors where you do not already have a SOC 2 report to review. Question 1 confirms whether one exists; if it does, review the report instead. Where practical, ask for supporting documentation rather than accepting a yes/no answer alone. The right-hand column is what to look for, not one correct answer. Log answers on the Vendor DDQ tab, one copy of the tab per vendor.

NO.	QUESTION	WHAT TO LOOK FOR
ORGANIZATION		
1	Does your organization have a current SOC 2 report (Type 2 or Type 1)?	If yes, review that report on the SOC 2 Report Review sheet instead. A Type 1 covers design at a point in time and gives less assurance than a Type 2, so some of these questions may still be worth asking.
2	Do you maintain a formal information security policy, and how often is it reviewed?	A named owner and a review date. Annual review is common.
3	Is there a designated security owner or equivalent role?	Yes. A named person accountable for the security program.
4	Do you train people who can access our data on security, and do you keep completion records?	Yes for people with access. Completion records are the usual evidence.
DATA HANDLING		
5	What categories of our data will you store or process?	The list matches the service you are buying. No undisclosed categories.
6	Where is our data stored (country or region)?	The location matches your contract and any regulatory commitment you have made.
7	How do you protect data at rest, including the algorithm and where keys live?	Teams often use AES-256 or another strong algorithm, with keys in a managed store. Ask what they actually use.
8	How do you protect data in transit?	TLS 1.2 or later on public endpoints is common. Ask how transit is protected.
9	How long is our data retained after the contract ends, and how is it deleted?	A defined retention period and a deletion method you can confirm.

NO.	QUESTION	WHAT TO LOOK FOR
10	Do you use our data to train models or for any purpose beyond the service?	Any secondary use should be disclosed. A no, or a contractual prohibition, is the usual expectation.
ACCESS CONTROL		
11	Do you enforce MFA for people who access systems that store or process our data?	MFA enforced for those systems is a common control. Ask whether it is required or optional.
12	How do you remove access when someone leaves or changes role?	A defined removal time. Same day or next business day is a typical commitment.
13	Is access to our data limited to people who need it for their role?	Yes. Least privilege, without standing admin access for people who do not need it.
14	Do you review access on a defined cycle?	Yes, on a set cycle. Ask for the last review.
INCIDENT RESPONSE		
15	Do you have a documented incident response plan, and have you exercised it?	A written plan with owners, and a record of an exercise or a real incident.
16	What is your commitment for notifying customers of a confirmed security incident?	State the contractual window. 72 hours is common; no trust services criterion sets it.
17	Have you had a security incident in the past 24 months that affected customer data?	Any yes should come with an explanation of impact and what changed afterward.
18	Do you keep a log of security incidents and near-misses?	A log is a common way to show the program is operating.
SUBPROCESSORS		
19	Do you share our data with subprocessors? If so, which ones, and what data?	A list, the data categories, and the purpose.
20	How do you assess a subprocessor's security before you use them?	A described vendor assessment, not an informal signup.
21	Will you notify us before adding or replacing a subprocessor that handles our data?	Yes, with a notice period you can live with.
BUSINESS CONTINUITY		
22	Do you have a documented business continuity and disaster recovery plan?	Yes, if they commit to continuity, and some sign it has been tested.

NO.	QUESTION	WHAT TO LOOK FOR
23	What recovery time and recovery point do you commit to?	Defined targets, if they are part of the service commitment. They should match the contract.
24	How often are backups taken, and how often do you test a restore?	A stated backup cadence and evidence restores are tested. Quarterly restore tests are common.
VULNERABILITY MANAGEMENT		
25	Do you test your application for vulnerabilities, including with a third party?	Scanning, remediation timing, and any third-party test. An annual pentest is common, not required.
26	How do you handle known vulnerabilities in your software dependencies?	A way to learn about them and a remediation timeline tied to severity.
27	Do you publish a vulnerability disclosure path or a security contact?	A monitored security contact. A formal disclosure program is typical for product companies.
LEGAL / COMPLIANCE		
28	Will you sign a data processing agreement?	Whether they will sign yours or offer their own. A DPA is typical when personal data is processed.
29	Do you carry cyber liability insurance? If so, what limit?	Some vendors carry it to transfer financial impact. It is not required evidence. Ask whether they do, and the limit if yes.
30	Which other frameworks apply to this service (for example HIPAA, PCI DSS, GDPR)?	Disclose the ones that actually apply. Do not treat a logo as the answer.
31	Will you support a periodic refresh of this review?	Yes, on a cycle you set.
TERMINATION		
32	How do you return or destroy our data when the contract ends?	A written process and a way they confirm deletion.
33	How long after contract end do you complete deletion?	A stated window. 30 to 90 days is a common contract term.
34	Can we export our data before the contract ends?	Yes, in a format you can actually use.
35	Do you have a documented exit or transition approach?	Yes if you will need help moving off the service. Ask what is included.

Need the Penetration Test?

The penetration test is one of the familiar pieces of a SOC 2 security program. For the audit itself, the question is often fairly straightforward: was the test performed, and is there evidence of it?

The next person who sees that report may be less forgiving.

An enterprise prospect's procurement or security team may ask how the application was tested, what methodology was used, what was in scope, how findings were rated, and whether identified issues were actually remediated. A report that was perfectly adequate for checking the SOC 2 box can become a problem when someone starts asking tougher questions.

That's where Asteros comes in.

We perform penetration tests the way we think they should be done: manual-first, methodical, and useful to the people who actually have to fix the application.

Our web application assessments are based on the OWASP Application Security Verification Standard (ASVS), with hands-on testing rather than a vulnerability scanner packaged as a pentest. Findings include reproduction steps, evidence, severity, and remediation guidance written for the technology being tested.

The ASVS assessment also documents the areas that were evaluated and what the application is doing well, so a clean result doesn't leave you wondering whether any work was really done.

One Report. Three Audiences.

The report is written for three audiences:

- ♦ **Executives** get a clear summary of what was found and what matters.
- ♦ **Engineers** get the technical detail needed to reproduce and fix issues.
- ♦ **Auditors and security reviewers** get the scope, methodology, severity documentation, and evidence behind the assessment.

After remediation, Asteros retests the findings at no additional charge, giving you documented evidence that identified issues were fixed and independently validated.

The goal isn't just to produce a pentest report.

It's to produce one that

What Asteros Does

METHOD	Manual-first web application penetration testing based on OWASP ASVS.
REPORTING	One report written for executives, engineers, and auditors/security reviewers.
RETEST	Remediation validation included at no additional charge.
PRICING	Flat-rate pricing with clear scope and deliverables.
WHO YOU WORK WITH	The senior tester doing the work, from scoping through reporting.

If this guide helped you get a better handle on SOC 2, that's the kind of useful, straightforward work we bring to the pentest itself.

Let's talk about your application.

asteros.com team@asteros.com

ASTEROS

© 2026 ASTEROS LLC • FREE TO USE AND SHARE