Kyūdō

SOVEREIGNTY-GRADE AI · GRC

SOC 2 Type II Preparation Guide

Vigilance with Purpose. Security with Control.

A practitioner's framework guide for organizations preparing their first or next SOC 2 Type II examination.

PRESENTED BY

Kyūdō — a KMicro Technologies platform

3525-265 Hyland Avenue

Costa Mesa, CA 92626 · kyudo.ai

hello@kyudo.ai

kyudo.ai · April 2026

CLASSIFICATION: PUBLIC

00 · HOW TO USE THIS GUIDE

Read this first.

This guide is written for the people who will actually carry the SOC 2 Type II examination on their backs: the security leader who owns the controls, the compliance lead who owns the evidence, the engineering manager who owns the change-management trail, and the executive who has to sign the management assertion. It is structured to be read end-to-end the first time and consulted in pieces every time after that.

Section 1 establishes what a SOC 2 Type II report actually is and is not, including the parts most teams get wrong before they even pick an auditor. Section 2 walks through the AICPA Trust Services Criteria with the level of specificity an auditor will use. Section 3 maps the Common Criteria — CC1 through CC9 — to the Microsoft Security stack most regulated mid-market organizations are already running, naming the exact signal source for each control. Section 4 covers the evidence model: what auditors expect, how sampling works, and what makes an artifact survive scrutiny. Section 5 is the twelve-month preparation timeline, week by week. Section 6 is the catalog of findings that show up most often and how to prevent them. Section 7 is the Kyūdō continuous-readiness model — not because every organization needs Kyūdō, but because the architectural pattern it describes is the direction every mature SOC 2 program eventually moves toward.

If you read nothing else, read Section 5. The single largest cause of a delayed or qualified Type II report is not a control failure. It is an evidence gap created by inconsistent operation during the observation window.

> **This guide is not legal or accounting advice**
> SOC 2 examinations are performed by licensed CPA firms under the AICPA's attestation standards. This guide is a practitioner reference that reflects the AICPA 2017 Trust Services Criteria with revised points of focus (2022) and current audit practice as of April 2026. Engage a qualified service auditor to scope, plan, and execute your examination.

01 · FOUNDATIONS

What SOC 2 Type II is, what it is not, and why the distinction matters.

The attestation, in one paragraph

SOC 2 is an attestation report issued by a licensed CPA firm under AICPA attestation standards. It opines on the design and operating effectiveness of a service organization's controls relevant to one or more of five Trust Services Criteria categories: Security, Availability, Processing Integrity, Confidentiality, and Privacy. Security — also called the Common Criteria — is required in every SOC 2 examination. The other four are optional and are added when the organization's service commitments and customer expectations require them.

Type I vs. Type II — the only difference that matters

A SOC 2 Type I report opines on whether controls are suitably designed and implemented at a single point in time — typically the report date. A SOC 2 Type II report opines on whether those same controls operated effectively over a period of time — the observation window, also called the audit period or examination period. Type I tests existence; Type II tests endurance.

Enterprise procurement teams have largely moved past Type I as a gating artifact. A Type I report is acceptable as a one-time bridge for a first-year service organization, but customers, partners, and insurers expect a Type II within twelve months. Treat Type I as a milestone, not a destination.

Dimension	SOC 2 Type I	SOC 2 Type II
Auditor opinion	Are controls suitably designed and implemented as of a specific date?	Did controls operate effectively across an observation period (typically 3, 6, or 12 months)?
Evidence required	Point-in-time artifacts: policies, configurations, screenshots.	Continuous evidence across the entire observation period: logs, tickets, completed reviews, sampled events.
Auditor procedures	Inquiry, inspection of design.	Inquiry, inspection, observation, and re-performance with sampling.
Typical timeline	3–6 months from kickoff to report.	6–12 months from kickoff to report (including the observation window itself).

Dimension	SOC 2 Type I	SOC 2 Type II
Renewal cadence	Treated as a one-time bridge; not renewed.	Annual; reports are conventionally treated as valid for 12 months.
Buyer acceptability	Tolerated for first-year vendors only.	Required by enterprise procurement, large healthcare systems, financial services buyers, and most cyber insurance underwriters.

Sources — AICPA, 2017 Trust Services Criteria with Revised Points of Focus (2022); AT-C Section 205, Examination Engagements.

Choosing the observation window

The observation window is the period across which the auditor will test operating effectiveness. The window length is a judgment call between three competing constraints: how soon you need a report in front of buyers, how much continuous evidence you can demonstrate, and how much auditor scrutiny the resulting report needs to withstand.

Window length	When to choose it	What the resulting report says about you
3 months	First-time Type II; you need a report in front of buyers urgently; you have just transitioned from a Type I and have established controls but limited operational history.	Acceptable as a bridge. Sophisticated buyers will note the short window and ask when the next 12-month report will be issued.
6 months	Recommended for first-time Type II. You have stable controls, you can demonstrate consistent operation, and you want a report that holds up in enterprise procurement.	The standard first-time Type II posture. Most procurement teams accept a 6-month report without comment.
12 months	Renewal cadence. Mature program. Annual examination aligned to fiscal year or buyer contract renewal cycle.	The industry-standard expectation for a mature service organization. A 12-month gap-free report is what enterprise buyers want to see.

After the first Type II report, the goal is a continuous twelve-month observation window with no gap between reports. A gap-free succession of 12-month reports is what large enterprise buyers, insurers, and regulators rely on. Any gap, even a short one, creates a period the auditor cannot opine on — and that is a question the buyer will ask.

The five categories, in plain terms

Security is required in every SOC 2 examination. The four trust categories beyond Security are added based on the organization's service commitments to its customers.

Trust category	What it covers	Add it when…
Security (Common Criteria)	Logical and physical access, system operations, change management, risk identification, incident response, governance and oversight. Mapped to COSO 2013 internal-control framework.	Always. Required for every SOC 2 examination.
Availability	Capacity planning, performance monitoring, environmental protections, backup and recovery, business continuity, disaster-recovery testing.	You commit to uptime, SLA, RTO, or RPO levels in customer contracts.
Processing Integrity	Inputs are valid, complete, accurate, and timely; processing produces output that meets the entity's specifications.	You process customer transactions, financial data, or workflows where output correctness is the contracted service.
Confidentiality	Data classified as confidential is identified, protected, retained, and disposed of according to commitments. Distinct from Privacy: applies to non-personal confidential data such as IP, trade secrets, or contract terms.	You handle confidential customer data, source code, financial records, M&A artifacts, or classified information under contractual confidentiality obligations.
Privacy	Personal information is collected, used, retained, disclosed, and disposed of in accordance with the organization's privacy commitments and the AICPA's Generally Accepted Privacy Principles. Aligns with GDPR, CCPA, and similar regimes.	You collect, store, use, or disclose personal information about identified individuals — especially in healthcare, financial services, education, or consumer-facing applications.

Source — AICPA TSP Section 100, 2017 Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy (with Revised Points of Focus — 2022).

A common error is to add too many categories in the first examination. Start with what your service commitments require. Adding categories later is a normal scope expansion. Withdrawing a category from a subsequent report is read by buyers as a regression.

02 · THE TRUST SERVICES CRITERIA

The framework auditors actually test against.

The 2017 Trust Services Criteria, with revised points of focus published in 2022, define what a service auditor evaluates. The criteria themselves are stable; the points of focus — the implementation guidance the AICPA provides for each criterion — were updated in 2022 to reflect modern threats, technologies, vulnerabilities, and the elevated expectations of enterprise buyers. Through 2025 and into 2026, the AICPA has not issued a wholly new version of the criteria. Practice has evolved through clarifying guidance and refined points of focus rather than a structural rewrite.

Two implications follow. First, the controls your organization is expected to have in place have grown more comprehensive: identity and access management, multi-factor authentication, encryption at rest and in transit, change management, logging, monitoring, incident response, vendor risk management, and — increasingly — governance over AI systems and the data that trains them. Second, auditors expect more substantive disclosure in the system description about how management identifies risk, how controls are designed to address those risks, and how controls operate over time.

Common Criteria structure

The Common Criteria are organized as nine series, CC1 through CC9, each of which corresponds to a category in the COSO 2013 internal-control framework. This structure is intentional: it lets a SOC 2 examination function as evidence of internal-control design and operating effectiveness in alignment with the framework most large enterprises and regulators already use.

Series	Title	What this series tests
CC1	Control Environment	Tone at the top, board oversight, organizational structure, accountability, and ethical commitment. The cultural and structural foundation everything else rests on.
CC2	Communication and Information	Internal and external communication of policies, responsibilities, and information needed to operate the control environment.
CC3	Risk Assessment	How the entity identifies, analyzes, and responds to risks to achieving its objectives, including risks of fraud and risks from change.
CC4	Monitoring Activities	Ongoing and separate evaluations of internal control performance and the communication of deficiencies.

Series	Title	What this series tests
CC5	Control Activities	The general policies, procedures, and technology controls that mitigate identified risks.
CC6	Logical and Physical Access Controls	Identity, authentication, authorization, segregation of duties, removal of access, and physical access to facilities and infrastructure.
CC7	System Operations	Detection of vulnerabilities and security events, incident response, recovery, and the operational backbone of the service.
CC8	Change Management	Authorization, design, implementation, and testing of changes to infrastructure, software, and data.
CC9	Risk Mitigation	Risk transfer (e.g. insurance, contractual provisions) and vendor and business-partner risk management.

Source — AICPA TSP Section 100, COSO 2013 Internal Control — Integrated Framework alignment.

The CC1 through CC5 series correspond directly to the five COSO 2013 components. CC6 through CC9 extend the framework to address logical and physical access, system operations, change management, and risk mitigation — the technical and operational controls most relevant to a software-driven service organization.

Points of focus — what they are and what they are not

Each criterion is supported by points of focus: enumerated examples of how an organization might implement the criterion. The 2022 revision expanded these in several areas, particularly around risk assessment, vendor risk, third-party reliance, and the governance of new technologies including AI.

Points of focus are illustrative, not mandatory. The AICPA states explicitly that not every point of focus must be addressed for a successful examination. The auditor's judgment is whether the controls in place, taken together, meet the criterion. That said, points of focus that are widely treated as table-stakes — multi-factor authentication for privileged access, formal vendor risk programs, encryption at rest and in transit, documented incident response — will functionally be expected. Treat them as obligations even though the framework calls them suggestions.

> **Where the 2022 revisions tightened expectations**
>
> Risk assessment must be a documented, recurring process, not a one-time exercise. Auditors expect to see the methodology, the people involved, the inputs, and the outputs over time.

Vendor risk management must extend beyond questionnaires to include continuous monitoring of subservice organizations and evidence that issues identified are tracked to closure.

Identity and access reviews must address privileged access, conditional access, and the prompt removal of access on termination or role change.

Change management must distinguish standard, normal, and emergency changes and provide evidence of authorization, testing, and approval for each.

AI governance is increasingly cited under risk assessment and processing integrity. Organizations using AI in customer-facing services should document model inventory, training-data provenance, and human-in-the-loop oversight before the auditor asks.

From criterion to control to native signal.

Most regulated mid-market organizations are already running the Microsoft Security stack — Microsoft Entra ID, Microsoft Defender XDR, Microsoft Defender for Cloud, Microsoft Sentinel, Microsoft Purview, and Azure Policy. These platforms generate the operational signal most SOC 2 controls require. The work is connecting the signal to the criterion. This section names the specific Microsoft service that produces the evidence for each Common Criteria series.

This is not exhaustive. It is the minimum viable mapping a service organization on the Microsoft stack should be able to demonstrate. Where your environment includes additional sources — a third-party MDM, a code-scanning platform, a separate ticketing system — those become additional evidence sources for the same controls.

CC1 — Control Environment

What it tests

Tone at the top, board oversight, organizational structure, accountability, ethical commitment, and the assignment of authority and responsibility for the design, implementation, and operation of the control environment.

Representative controls

- Documented and approved code of conduct or acceptable-use policy, signed by all personnel.
- Defined organizational structure with clear reporting lines for security, compliance, and IT.
- Board or executive oversight of the security and compliance program, with documented review cadence.
- Background checks for personnel in security-sensitive roles.
- Annual security and privacy training, with completion tracked and remediated for non-completers.

Microsoft and adjacent signal sources

Evidence	Where it lives
Policy acknowledgment records	HRIS, learning management system, or Microsoft Forms responses logged to SharePoint/Teams.
Training completion records	LMS reporting; for Microsoft-native training, Microsoft Defender for Office 365 attack-simulation completion reports and Viva Learning records.
Org chart and role assignments	Microsoft Entra ID groups, dynamic group membership rules, and administrative-unit assignments.

Evidence	Where it lives
Board reporting	Documented meeting minutes, agendas, and packets stored in a controlled SharePoint or Teams location with retention enforced via Microsoft Purview.

CC2 — Communication and Information

What it tests

How relevant information is generated, captured, and used to support the operation of the control environment, and how that information is communicated internally and externally.

Representative controls

- Approved policy and procedure library, version-controlled and accessible to personnel.
- Documented service commitments, system requirements, and customer-facing security commitments.
- Internal communication channels for reporting suspected security or compliance issues, including an anonymous channel.
- Customer communication procedures for system status, incidents, and material changes to commitments.

Microsoft and adjacent signal sources

Evidence	Where it lives
Versioned policy library	SharePoint with versioning enabled; Microsoft Purview retention labels enforce policy lifecycle.
Customer-facing commitments	Master Service Agreement, Data Processing Addendum, and Trust Center pages, archived with timestamps.
Internal incident reporting channel	Microsoft Teams compliance channel or dedicated mailbox; anonymous reporting via third-party tool integrated to Teams.
Status communication	Customer status page, Microsoft 365 admin alerts, and tenant-level email distribution to designated customer contacts.

CC3 — Risk Assessment

What it tests

How the organization identifies and analyzes risks to the achievement of its objectives, considers fraud risk, and identifies and assesses changes that could significantly impact the system of

internal control. The 2022 revisions sharpened this area: auditors expect a documented, recurring process.

Representative controls

- Annual enterprise risk assessment, with a documented methodology covering threat, vulnerability, likelihood, and impact.
- Risk register that is reviewed at a defined cadence and updated for new threats, technologies, and business changes.
- Documented risk acceptance, mitigation, transfer, and avoidance decisions, signed by an accountable owner.
- Fraud risk assessment that addresses both internal and external fraud scenarios.
- Change-driven risk assessment triggered by material changes (new product, new region, new processor, new AI model in customer-facing use).

Microsoft and adjacent signal sources

Evidence	Where it lives
Risk register and assessment artifacts	Centralized in a controlled SharePoint or GRC platform location, with access restricted via Microsoft Entra ID groups.
Vulnerability identification feed	Microsoft Defender Vulnerability Management, Defender for Cloud recommendations, and third-party scanners exporting to Sentinel.
Threat intelligence input	Microsoft Defender Threat Intelligence, Sentinel threat-intelligence connectors, ISAC/ISAO feeds.
Change-driven assessment evidence	Architecture review board minutes, change advisory board records, AI Governance Board minutes (where applicable).

CC4 — Monitoring Activities

What it tests

Ongoing evaluations, separate evaluations, or some combination thereof, used to ascertain whether the components of internal control are present and functioning. Communication of internal control deficiencies in a timely manner to those parties responsible for taking corrective action.

Representative controls

- Continuous monitoring of security controls (SIEM detections, control-state telemetry, configuration drift).
- Periodic internal control assessments, with findings tracked to remediation.

- Independent assessments — internal audit, third-party penetration testing, vulnerability scans — conducted at a defined cadence.
- Defined process for escalating identified deficiencies to management and the board.

Microsoft and adjacent signal sources

Evidence	Where it lives
Continuous control monitoring	Microsoft Sentinel analytic rules, Microsoft Defender for Cloud Secure Score over time, Microsoft Purview Compliance Manager assessment scores.
Internal assessments	Internal audit reports, control self-assessments, and remediation tickets tracked in Azure DevOps or Jira.
External assessments	Annual penetration test report, external vulnerability scan reports, prior-year SOC 2 management responses.
Deficiency escalation	Documented escalation matrix and evidence of escalations occurring (e.g. tickets, email chains, board packet excerpts).

CC5 — Control Activities

What it tests

Selection and development of control activities that contribute to the mitigation of risks to the achievement of objectives, including general controls over technology and policies that establish what is expected and procedures that put policies into action.

Representative controls

- Documented control matrix mapping each control to the criterion it addresses, the owner, the operating frequency, and the evidence.
- Segregation of duties for security-sensitive operations (production access, financial transactions, change deployment).
- Documented standard operating procedures for security operations, change management, and incident response.
- Periodic review of control design effectiveness, particularly after material changes to systems or processes.

Microsoft and adjacent signal sources

Evidence	Where it lives
Control matrix	Maintained in the GRC platform of record; exported on request as a control-by-control evidence index.

Evidence	Where it lives
Segregation of duties	Microsoft Entra ID Privileged Identity Management role assignments; conflicting-role detections and reviews.
Procedures	SharePoint or Teams document library, version-controlled, with named owners and review dates.
Design review evidence	Architecture review minutes, change advisory board records.

CC6 — Logical and Physical Access Controls

What it tests

This series is the largest by control count in most SOC 2 reports. It tests how the organization restricts logical access, prevents or detects unauthorized access, manages identification and authentication, manages credential lifecycles, addresses physical access, manages the disposal of physical media, and protects information during transmission, movement, and removal.

Representative controls

- Multi-factor authentication enforced for all interactive logins, with enhanced verification for privileged actions.
- Role-based access control with documented role definitions and least-privilege assignment.
- Quarterly access reviews of all in-scope systems, with documented attestation by control owners.
- Just-in-time elevation for privileged access; standing administrative privileges are minimized and time-bounded.
- Documented joiner-mover-leaver process with same-day access removal for terminations.
- Encryption at rest and in transit for all customer data; documented key management procedures.
- Physical access controls for data centers (typically inherited from cloud subservice organizations) and corporate facilities.
- Secure disposal procedures for physical media.

Microsoft and adjacent signal sources

Evidence	Where it lives
MFA enforcement	Microsoft Entra Conditional Access policies; Entra ID sign-in logs filtered for MFA satisfaction; Authentication Methods activity report.
Privileged access	Microsoft Entra Privileged Identity Management activation logs, role assignment audit logs, JIT activation records, and access reviews.
Access reviews	Microsoft Entra ID Access Reviews — review records, attester names, decisions, and remediation actions.
Joiner-mover-leaver	Entra ID provisioning logs, HRIS-to-Entra provisioning audit, group membership changes, and license assignment changes.
Encryption at rest	Azure Storage encryption settings, Azure SQL Transparent Data Encryption status, customer-managed keys in Azure Key Vault.

Evidence	Where it lives
Encryption in transit	TLS configuration evidence (Azure Front Door, Application Gateway, App Service settings), certificate inventory in Azure Key Vault.
Physical access — data center	Inherited from Microsoft Azure subservice organization; covered by the Azure SOC 2 report (a key complementary subservice control).
Conditional access posture	Conditional Access policy export, what-if simulation logs, sign-in logs showing CA decisions.

CC7 — System Operations

What it tests

Detection of vulnerabilities and security events, response to incidents, and the recovery of affected systems. This is where the SOC, the SIEM, the incident response runbook, and the business-continuity program are tested as a single connected capability.

Representative controls

- Continuous vulnerability identification on infrastructure, applications, and endpoints; documented remediation SLAs by severity.
- Centralized logging and monitoring with documented retention; correlation rules for security events of interest.
- 24x7 incident detection and response capability, with named on-call roles and escalation paths.
- Incident response procedure tested at least annually (tabletop or live exercise) with lessons learned documented.
- Documented and tested business continuity and disaster recovery procedures, with RTO and RPO aligned to customer commitments.
- Customer breach notification procedure consistent with regulatory and contractual obligations.

Microsoft and adjacent signal sources

Evidence	Where it lives
Vulnerability identification	Microsoft Defender Vulnerability Management findings and remediation history; Defender for Cloud recommendations and resolved status.
SIEM and detection	Microsoft Sentinel — analytic rule definitions, incident records, MTTR metrics, hunting query history.

Evidence	Where it lives
Endpoint protection	Microsoft Defender for Endpoint posture, alert handling history, automated investigation outcomes.
Identity threat detection	Microsoft Defender for Identity, Entra ID Identity Protection risk detections and remediations.
Cloud workload protection	Microsoft Defender for Cloud workload protection alerts and resolutions.
Incident records	Sentinel incidents — status, ownership, evidence, root-cause analysis attached to each incident.
Tabletop exercises	Documented exercise scenarios, participant lists, after-action reports, and remediation tickets.
Backup and recovery	Azure Backup recovery point success rate, recovery test reports, restore-time measurements against RTO.
DR testing	Azure Site Recovery test failover reports; DR runbook execution evidence.

CC8 — Change Management

What it tests

Authorization, design, development, configuration, documentation, testing, approval, and implementation of changes to infrastructure, data, software, and procedures. The 2022 revisions place particular weight on distinguishing change types and producing evidence of approval and testing for each.

Representative controls

- Documented change-management policy with defined change types (standard, normal, emergency) and approval requirements.
- Pre-production testing for all material changes; documented test results.
- Code review enforced via branch protection or equivalent; commits to production branches require peer approval.
- Production-deployment authorization is logged and reconstructable.
- Configuration baseline management for production infrastructure; drift detection.
- Emergency-change procedure with after-the-fact documentation and review.

Microsoft and adjacent signal sources

Evidence	Where it lives
Change records	Azure DevOps Boards, Jira, or ServiceNow change

Evidence	Where it lives
	tickets — type, approver, test evidence, deployment outcome.
Code review enforcement	Branch protection rules in Azure Repos, GitHub, or GitLab; pull-request approval history.
Pipeline evidence	Azure DevOps Pipelines or GitHub Actions run history with stage approvals, test results, and deployment artifacts.
Configuration baseline and drift	Azure Policy compliance state over time; Azure Resource Graph queries; Defender for Cloud regulatory compliance.
Emergency-change records	Documented emergency-change tickets with retroactive review and CAB sign-off.

CC9 — Risk Mitigation

What it tests

How the entity identifies, selects, and develops risk mitigation activities for risks arising from potential business disruptions, and how it assesses and manages risks associated with vendors and business partners. This is where vendor risk management is tested in earnest.

Representative controls

- Documented vendor risk management program with tiered classification by criticality and data sensitivity.
- Pre-engagement security review for all material vendors, including evidence review (SOC 2, ISO 27001, security questionnaires).
- Continuous monitoring of subservice organizations, with documented response when their security posture changes.
- Cyber insurance coverage aligned to risk profile; documented review of coverage adequacy.
- Documented response procedures for subservice-organization incidents that may affect customer data.

Microsoft and adjacent signal sources

Evidence	Where it lives
Vendor inventory	Maintained in the GRC platform of record; reconciled against accounts payable and SaaS-discovery sources.
Vendor security reviews	Stored evidence: SOC 2 reports, security questionnaires, contracts with security and breach-

Evidence	Where it lives
	notification clauses.
Subservice monitoring	Subscriptions to vendor trust centers, RSS feeds, and security advisories; documented response actions.
Insurance evidence	Cyber-insurance certificate of coverage, coverage limits, and renewal history.

04 · EVIDENCE

What auditors actually accept — and what they reject.

The four procedures auditors use

In a Type II examination, auditors apply four procedures to test operating effectiveness. Strong evidence is produced when artifacts align with all four, not just one.

Procedure	What it answers	What it produces as evidence
Inquiry	How is the control performed?	Auditor interview notes; control owner walkthroughs.
Inspection	Does the artifact exist as expected?	Reviewed policies, configurations, screenshots, exported reports, signed approvals.
Observation	Is the control performed in real time as described?	Auditor observation of a live control execution — a deployment, an access review, an incident triage.
Re-performance	Does the control work when the auditor exercises it?	Auditor re-runs the control or recreates the result independently — e.g. recompiles an access list, reproduces an alert.

Source — AT-C Section 205, Examination Engagements; auditor practice guides.

A control supported only by a screenshot is supported only by inspection. A control supported by a screenshot, a policy, an observed execution, and a re-performance is supported by all four. Type II evidence quality is the difference between a clean opinion and a list of management responses.

Sampling — how the auditor decides what to test

Type II observation periods cover thousands of control executions. No auditor reviews all of them. Sampling is professional judgment guided by the population, the risk, the expected deviation rate, and the tolerable deviation rate. In practice, auditors sample to achieve representative coverage across the period — they will deliberately pick events from early, middle, and late in the window, and will deliberately include high-risk periods (releases, incidents, leadership changes).

Control frequency	Typical sample expectation	What this means for you
Daily	25–40 instances across the period.	Automated controls and continuous evidence are essential. Manual daily controls are very expensive to evidence at this sample size.
Weekly	10–15 instances.	Sample selection will frequently target weeks containing releases, incidents, or holidays.
Monthly	All months sampled, frequently every month inspected.	A single missed month creates a finding. A skipped month around year-end is the single most common Type II finding.
Quarterly	All quarters inspected.	Quarterly access reviews are routinely sampled in full. The auditor will check both that the review occurred and that the remediation actions identified were closed.
Annual	The single instance in the period is inspected.	Annual controls (penetration test, BCP exercise, risk assessment, training) are tested against a single artifact — which means that one artifact must be flawless.

Sampling expectations vary by auditor, control risk, and population characteristics. Confirm the expected sample plan with your service auditor before fieldwork begins.

What makes evidence audit-defensible

Auditors are looking for artifacts that answer five questions without follow-up. Strong evidence answers all five on its face.

Question	What audit-defensible evidence shows
Who?	Named, identifiable actor — a specific user, system, or service principal. Generic accounts and shared credentials are not acceptable.
When?	Timestamp from a trusted source. The artifact's metadata, the system log, or the email header must establish the time.
What?	The action performed, in terms specific enough to map directly to the control. ‘Access review completed’ is too vague; ‘Quarterly Tier-1 administrator access review for production tenant ABC, completed and certified by [name] on [date]’ is acceptable.

Question	What audit-defensible evidence shows
Why?	The control or policy that authorizes the action, referenced by its identifier. The artifact connects back to the control matrix.
How is it tamper-evident?	Source-system logs from a system the actor cannot edit; cryptographic hashes; chain of custody. A screenshot edited in Photoshop is indistinguishable from a real one. A SIEM-exported log with a query hash is not.

The screenshot problem

Screenshots are the most common SOC 2 evidence and the weakest. They show a moment, not a duration. They can be edited. They cannot be re-performed. The 2022 revisions and current auditor practice push toward direct system extracts, API exports, and timestamped logs as the primary evidence form.

Where a screenshot is genuinely the best available artifact — a vendor portal that does not expose an API, a configuration screen that cannot be exported — it must be paired with metadata that establishes when it was taken, by whom, and from what system, and ideally with a hash anchored to a tamper-evident store.

The continuity test

The single most common Type II finding is not a missing control. It is an inconsistent control: an access review performed in months 1, 2, and 3, then skipped in month 4, then resumed. A vulnerability scan run weekly, then biweekly during the holidays. A quarterly review where the third quarter's evidence cannot be located.

Auditors call this a deviation. A pattern of deviations — even small ones — produces a qualified opinion or a list of management responses appended to the report. The remedy is not heroic effort during fieldwork. The remedy is automation and discipline during the observation window.

05 · PREPARATION TIMELINE

Twelve months from kickoff to clean opinion.

This timeline assumes a six-month observation window for a first-time SOC 2 Type II report and three months of preparation before the window opens, plus three months of fieldwork and reporting after it closes. Adjust the timeline to your chosen window length: a three-month window compresses Phase 3; a twelve-month window expands it. The phases themselves do not change.

Phase 1 — Scoping and design (months 1–2)

Goal

Determine which Trust Services Criteria categories apply, draft the system description, identify the controls that satisfy each criterion, and confirm gaps.

Activities

1. Define the scope: which products, services, environments, and customer commitments are in scope. The scope decision determines everything else.
2. Select the trust categories beyond Security based on customer commitments. Document the rationale.
3. Engage a service auditor. The auditor must be a licensed CPA firm. Negotiate the engagement letter, the observation window, and the fieldwork schedule.
4. Draft the system description: organization overview, services provided, principal service commitments, system requirements, infrastructure, software, people, procedures, and data.
5. Build the control matrix: each criterion, the control that addresses it, the owner, the frequency, the evidence source.
6. Run a readiness assessment — internal or with an advisory firm — to identify gaps before the observation window opens.

Exit criteria

- Signed engagement letter with the service auditor.
- Approved scope, system description draft, and control matrix.
- Readiness assessment findings logged with named remediation owners and target close dates.

Phase 2 — Remediation (months 2–3)

Goal

Close every readiness-assessment gap before the observation window opens. Controls that do not exist on day one of the observation window cannot be evidenced — and a missing control cannot be backdated.

Activities

1. Implement missing controls; do not document controls that do not yet operate.
2. Stand up evidence collection at source. Configure log forwarding, ticketing-system tagging, access-review automation, and SIEM analytics so that controls produce evidence as a byproduct of operating.
3. Establish the evidence repository structure: a folder, library, or GRC module per control, with consistent naming and date stamping.
4. Train control owners on the evidence expectations for their controls. The first time an owner sees the requirement should not be when fieldwork starts.

Exit criteria

- All controls in the matrix are operating.
- Evidence collection is automated wherever possible and scheduled wherever it is not.
- The evidence repository is structured, named, and access-controlled.
- Control owners can produce evidence on demand for their controls without auditor intervention.

Phase 3 — Observation window (months 3–9, six-month window)

Goal

Operate every control consistently. Capture evidence as a byproduct, not as an extraction event. Maintain the discipline that turns three good months into six gap-free months.

Activities

1. Execute every control on its operating frequency. No skips, no shortcuts, no ‘we will do it next month’.
2. Capture evidence at execution time — not at month-end. Auditors look for the timestamp on the artifact, not the date you uploaded it.
3. Run monthly evidence-quality reviews: confirm every control has produced the expected evidence for the month, and that the evidence answers who/when/what/why on its face.
4. Maintain a deviation log. When a control miss happens — and one usually does — record what happened, when it was discovered, the corrective action, and the prevention measure. Auditors do not penalize an honest deviation log; they penalize undisclosed deviations.

5. At the midpoint of the window, run an internal mock audit. Have someone outside the program try to find evidence for a sampled selection of controls. Track time-to-find and quality of evidence produced.

Exit criteria

- Every control has continuous evidence across the entire window.
- Deviation log is complete and each deviation has a documented corrective action.
- Mock audit results show evidence is producible on demand for any sampled control.

Phase 4 — Fieldwork (months 9–11)

Goal

Support the auditor's testing efficiently. Auditor questions answered quickly produce shorter fieldwork and fewer follow-up requests. Auditor questions that linger produce more requests, more deviations, and a longer engagement.

Activities

7. Designate a single audit liaison who triages all auditor requests and routes them to control owners. One throat to choke for both sides.
8. Stand up a shared evidence channel: a Teams channel, a SharePoint folder, or a GRC platform module that gives auditors structured access to evidence without ad-hoc emails.
9. Track every auditor request: who asked, what was asked, when it was answered, what was provided. Maintain a request log that becomes a record of the engagement.
10. Resolve exceptions promptly. When the auditor identifies a deviation, document the corrective action, the root cause, and the prevention measure for the next period.

Exit criteria

- All auditor requests answered and evidence delivered.
- Identified exceptions documented with management responses.
- Draft report received and reviewed by management.

Phase 5 — Reporting and renewal (months 11–12, then continuous)

Goal

Receive a clean opinion. Distribute the report. Set up the next observation window so there is no gap between reports.

Activities

5. Review the draft report carefully. Ensure the system description matches reality, the controls are described accurately, and any management responses to exceptions are clearly written.
6. Sign the management assertion and receive the final report.

7. Distribute the report through the controlled channels: customer trust portal, NDA-gated download, direct delivery to procurement contacts.
8. Schedule the next observation window. The next window should begin the day after the prior window ends. Any gap is a gap in continuous coverage.
9. Run a post-engagement retrospective. What controls produced evidence cleanly? Which ones required heroic effort? What automation needs to be in place before the next window opens?

Exit criteria

- Signed report distributed to customers and partners.
- Next observation window opens with no gap from the prior window's end.
- Retrospective findings logged into the program backlog.

06 · COMMON FINDINGS

Eight findings that show up in nearly every first-time examination.

Every Type II report has the potential for findings, and most first-time examinations have at least one. The patterns are remarkably consistent. The eight findings below account for a clear majority of qualified opinions and management responses in first-time SOC 2 Type II reports. Each is preventable. None are technical novelties.

1. Incomplete or inconsistent access reviews

The pattern: a quarterly access review is documented for the first quarter of the observation window, partially documented for the second, missing for the third, and rushed for the fourth. The remedy is to automate the review trigger, the reviewer assignment, and the attestation capture. In the Microsoft stack, Microsoft Entra ID Access Reviews provides this end-to-end. In a multi-platform environment, the same pattern is replicated for every platform with privileged access — cloud platforms, on-prem directories, SaaS applications with administrative roles.

Prevent it by scheduling the access review as a recurring item with a named reviewer, an automated reminder, and an automated escalation if the review is not completed by the due date. The output is a record — a list of accounts, the reviewer, the decision per account, and the timestamp — that survives sampling.

2. Stale or untested incident response procedures

The pattern: an incident response runbook exists but has not been exercised in the observation window, or has been exercised in name but produced no documentation. Auditors expect at least one tabletop exercise within the window with a documented scenario, a participant list, an after-action report, and remediation tickets for any issues identified.

Prevent it by scheduling the tabletop in advance, recording the session, and producing a structured after-action report. The exercise itself is the evidence. The after-action report is what the auditor will sample.

3. Vendor risk management that stops at onboarding

The pattern: vendors are assessed at onboarding with a security questionnaire and a SOC 2 review. After that, no recurring monitoring occurs. When a subservice organization has a security incident or material change, the service organization has no record of how it was identified, escalated, or responded to. The 2022 revisions sharpened expectations here significantly.

Prevent it by establishing a vendor monitoring cadence: subscriptions to vendor trust centers, automated alerts on vendor security advisories, an annual vendor reassessment for all vendors above a defined criticality tier, and a documented response procedure when a subservice organization changes its security posture or has an incident.

4. Change management without evidence of approval

The pattern: a change management policy exists, and changes are deployed via CI/CD pipelines, but the deployment evidence does not connect back to a documented approval. Auditors will sample deployments and ask: ‘Show me the change request, the approval, the testing, and the deployment for this change.’ If the answer involves multiple systems and a manual reconciliation, that is a finding.

Prevent it by tying every production deployment to a change record, with the approval and the testing evidence captured automatically by the pipeline. Branch protection rules in Azure Repos, GitHub, or GitLab provide this for code changes. Infrastructure-as-code changes need the same treatment with explicit ticket references in commit messages or pipeline metadata.

5. Gaps in privileged access monitoring

The pattern: privileged access is granted but not actively monitored. The organization can demonstrate who has privileged roles, but cannot demonstrate when privileged actions occur, who performs them, or whether anyone reviews them. Just-in-time elevation, where privileged access is requested, granted for a defined window, and logged, materially reduces this risk.

Prevent it by enforcing PIM (Privileged Identity Management) for all privileged roles. PIM activation logs become the evidence: who, when, why, and how long. Access reviews on privileged role assignments become the periodic control.

6. Encryption claims without key management evidence

The pattern: a policy states that data is encrypted at rest using AES-256, but the auditor cannot find evidence of key rotation, separation of duties on key access, or recovery procedures if a key is compromised. Encryption-at-rest claims must be paired with key-management evidence to be defensible.

Prevent it by managing keys in Azure Key Vault (or equivalent), enabling soft-delete and purge protection, configuring rotation policies, restricting access via Conditional Access and PIM, and producing periodic key-inventory and rotation reports as control evidence.

7. Backup procedures without recovery testing

The pattern: backups run successfully every night for the entire window. No recovery test is performed. The auditor's question is the obvious one: do the backups actually work? Recovery testing is a documented restore from backup to a test environment, with the result validated and recorded.

Prevent it by scheduling at least one recovery test within the observation window, ideally per quarter, with the scope, the result, and the time-to-recovery measured against the documented RTO. The artifact is the test report.

8. Logging gaps and SIEM blind spots

The pattern: critical systems are not logging to the SIEM, or logs are retained for less than the required duration, or alerts that should be tuned to the organization's environment are using

default vendor configurations. Auditors will check what is being logged and how long it is retained against the organization's stated commitments.

Prevent it by defining a logging architecture that explicitly enumerates the sources, the destination (typically Microsoft Sentinel for Microsoft-stack organizations), the retention period, and the analytic rules. Validate quarterly that all in-scope sources are still logging. New systems added during the window must be onboarded to the SIEM as part of the deployment process, not as a follow-up.

07 · THE KYŪDŌ MODEL

Continuous readiness as a condition, not a project.

Every section of this guide describes the same operating principle: the organizations that produce clean SOC 2 Type II reports without heroic effort are the organizations that operate their controls continuously and capture evidence as a byproduct. The audit cycle is uneventful because the readiness was already there. This is not a function of audit-prep tooling. It is a function of architecture.

This is the operating principle Kyūdō was built to express. The patterns that follow describe how a continuous-readiness program works at the architectural level. They are useful whether or not you ever evaluate Kyūdō specifically.

Architecture pattern — governance inside the security boundary

The conventional GRC architecture deploys a SaaS governance platform that ingests evidence from the organization's environment, processes it externally, and presents posture in the vendor's cloud. This pattern requires data egress: control-state telemetry, evidence artifacts, and sometimes raw logs leave the organization's environment to be governed.

The continuous-readiness pattern inverts this. The governance layer deploys inside the organization's own security boundary — in regulated organizations on the Microsoft stack, this means inside the customer's Azure tenant. Native integrations read from Microsoft Entra ID, Microsoft Defender, Microsoft Sentinel, Microsoft Purview, and Azure Policy directly, with no data egress. Evidence is produced and stored where the data already lives. Sovereignty is the architecture, not a configuration option.

> **Sovereignty as architecture**
> The Kyūdō platform deploys as microservices inside the customer's Azure tenant. Private endpoints, system-assigned managed identities, and customer-managed encryption keys mean no governance data crosses the tenant boundary. The deployment model is the moat — no SaaS-first competitor can replicate it without re-architecting their platform.

One control set, every framework

Most regulated organizations answer to multiple frameworks: SOC 2, ISO 27001, NIST CSF, CMMC, HIPAA, PCI DSS, GLBA, EU GDPR, and increasingly EU AI Act and ISO 42001. The conventional approach maintains a separate evidence pipeline per framework, which produces redundant work, inconsistent results, and the audit-cycle scramble this guide opens with.

The continuous-readiness pattern uses a meta-framework as the substrate — the Secure Controls Framework (SCF), with 1,470+ controls across 80+ frameworks. Test once, satisfy many. SOC 2's CC6.1 maps to ISO 27001:2022 Annex A.5.15, NIST CSF v2.0 PR.AA, CMMC AC.L2-3.1.1, HIPAA Security Rule §164.312(a)(1), and so on. The evidence that satisfies one criterion satisfies all related criteria across every framework the organization answers to.

Evidence that recalculates

The traditional GRC platform stores evidence as it is collected: a static artifact attached to a control. The continuous-readiness pattern recalculates control state on every signal. When Defender for Endpoint reports a new device, when Conditional Access blocks a sign-in, when Purview labels a new dataset, the relevant controls update their state in real time. This is what 'evidence that is already true between audits' means: the evidence does not need to be assembled because it was never disassembled.

Kyūdō's Capability Maturity & Completeness Assessment Engine (CMCAE) recalculates completeness, capability maturity levels, and residual risk on every signal and on every control change. The audit story does not need to be written. It is being maintained continuously.

Auditor-defensible AI

AI in GRC is now a category-saturated claim. Most platforms position AI as a chat layer over documents. The continuous-readiness pattern requires AI that survives the auditor's next question: every AI-produced explanation, mapping, or recommendation must have a source, a confidence level, and a re-performable result. Where AI participates in scoring or state transitions, those operations are deterministic functions; AI is advisory only.

This two-layer trust architecture — deterministic functions for scoring and state, AI for explanation and acceleration — is what makes the AI-produced output of a continuous-readiness platform admissible as control evidence. It is the difference between a chat-over-PDF tool and a system of record.

Where this leaves you

If you are preparing your first SOC 2 Type II examination, you do not need a continuous-readiness platform. You need scope, controls, evidence discipline, and a six-month observation window operated cleanly. This guide gives you the playbook for that.

If you are preparing your second or third Type II, and the audit-cycle scramble has gotten worse rather than better as your organization has grown, the architecture this section describes is the direction the practice is moving. The marginal cost of one more framework, one more region, one more product line should approach zero. If it does not, the bottleneck is the architecture.

Kyūdō is the platform that makes that architecture available to regulated organizations running Microsoft 365 and Azure. The next step, if useful, is a deployment workshop in your tenant. The architecture brief is one click. The conversation is one email.

—

If this is useful, the next step is concrete

Architecture briefing — a 30-minute walkthrough of the Kyūdō deployment in your Azure tenant: controls, evidence flow, and the sovereignty model. → hello@kyudo.ai

Controls workshop — 90 minutes mapping your current SOC 2 control set to the continuous-evidence model, with side-by-side framework views. → kyudo.ai/workshop

Trust packet — our SOC 2 posture, architecture diagram, data-residency statement, and the Microsoft estate dependency map. Available on request.

APPENDIX A · GLOSSARY

Terms used in this guide.

Term	Definition
AICPA	American Institute of Certified Public Accountants. The professional body that publishes the SOC 2 framework, the Trust Services Criteria, and the attestation standards under which SOC 2 examinations are performed.
Attestation	An engagement in which a CPA firm expresses a written conclusion about the reliability of subject matter against criteria. SOC 2 reports are attestation reports.
AT-C Section 205	The section of the AICPA's attestation standards that governs examination engagements. SOC 2 Type I and Type II are examination engagements.
Common Criteria	The Security category of the Trust Services Criteria, organized as nine series (CC1–CC9). Required in every SOC 2 examination.
COSO 2013	The Committee of Sponsoring Organizations Internal Control — Integrated Framework. The five COSO components map to the Common Criteria CC1–CC5.
Control matrix	The structured inventory of controls, their owners, their operating frequencies, the criteria they address, and the evidence that demonstrates their operation.
Deviation	A control execution that did not occur as designed during the observation window. Deviations are testable findings that may, depending on severity, result in a qualified opinion.
Examination period	The period across which the auditor tests operating effectiveness. Also called the observation period, audit period, monitoring period, or observation window.
Management assertion	The written statement by management asserting that the system description is accurate and the controls are suitably designed and (for Type II) operating effectively. The assertion is signed by an executive of the service organization.
Points of focus	Illustrative examples of how an organization might implement each Trust Services Criterion. Updated by the AICPA in 2022. Not mandatory, but functionally expected in many cases.

Term	Definition
Qualified opinion	An auditor opinion that contains exceptions — controls that did not operate effectively. A qualified opinion is acceptable in some circumstances but signals issues to enterprise buyers.
Re-performance	An audit procedure in which the auditor independently re-executes the control or recreates the result to verify operating effectiveness.
Scope	The boundary of the SOC 2 examination: the products, services, environments, and customer commitments that the report covers.
SCF	Secure Controls Framework. A meta-framework with 1,470+ controls across 80+ frameworks, used to test once and satisfy many.
Service auditor	The licensed CPA firm that performs the SOC 2 examination and issues the report.
Service organization	The entity whose controls are being examined. The organization that publishes the SOC 2 report to its customers.
Subservice organization	A vendor whose controls the service organization relies upon (e.g. a cloud provider, a payroll processor). May be addressed via the carve-out method or the inclusive method in the report.
System description	The narrative description of the service organization's system, including infrastructure, software, people, procedures, data, principal service commitments, and system requirements. Required by the SOC 2 Description Criteria.
Trust Services Criteria	The AICPA-published criteria against which SOC 2 examinations are performed. Five categories: Security, Availability, Processing Integrity, Confidentiality, Privacy.
TSP Section 100	The AICPA publication containing the Trust Services Criteria. Most recent version: 2017 with Revised Points of Focus, 2022.
Type I	A SOC 2 report on the suitability of design of controls at a specific point in time.
Type II	A SOC 2 report on the suitability of design and operating effectiveness of controls across an observation window.

APPENDIX B · AUTHORITATIVE SOURCES

Where to verify and go deeper.

This guide reflects the AICPA Trust Services Criteria and current audit practice as of April 2026. The primary sources below are authoritative; secondary sources are useful for practitioner perspective. Always defer to the AICPA publications and your engaged service auditor for definitive guidance.

Primary — AICPA

- AICPA TSP Section 100, 2017 Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy (with Revised Points of Focus — 2022). The authoritative criteria document.
- AICPA 2018 SOC 2 Description Criteria (with Revised Implementation Guidance — 2022). Defines what the system description must contain.
- AICPA AT-C Section 205, Examination Engagements. The attestation standard governing SOC 2 examinations.
- AICPA Guide, Reporting on an Examination of Controls at a Service Organization Relevant to Security, Availability, Processing Integrity, Confidentiality, or Privacy. The practitioner guide for service auditors.

Secondary — framework alignment

- COSO 2013 Internal Control — Integrated Framework. Maps to Common Criteria CC1–CC5.
- Secure Controls Framework (SCF) v2025.x. Cross-framework meta-control set with SOC 2 mappings.
- ISO/IEC 27001:2022 and ISO/IEC 27002:2022. Frequently cross-mapped to SOC 2 in SOC 2+ examinations.
- NIST Cybersecurity Framework v2.0. Frequently cross-mapped to SOC 2.

Microsoft documentation

- Microsoft Service Trust Portal — Microsoft's own SOC 2 report, available under NDA. Establishes the inherited subservice controls for Microsoft Azure and Microsoft 365 customers.
- Microsoft Purview Compliance Manager — SOC 2 assessment template and improvement actions mapped to Microsoft 365 and Azure controls.
- Microsoft Defender for Cloud — SOC 2 regulatory compliance dashboard and recommended controls.

Practitioner perspectives

- Selected articles from Linford & Co., BrightDefense, SecureFrame, Drata, Compyl, Konfirmity, and Strike Graph were consulted for current 2025–2026 audit practice patterns.

APPENDIX C · PRE-ENGAGEMENT CHECKLIST

Twenty-five items. Run them before you sign the engagement letter.

If you can answer ‘yes, with evidence’ to every item below, your first Type II examination is positioned for a clean opinion. If you cannot answer ‘yes’ to ten or more, do a readiness assessment first.

Scope and design

- The scope of the examination is documented and approved — specific products, services, environments, customer commitments.
- Trust categories beyond Security have been selected based on documented service commitments.
- The system description draft exists and addresses all required elements (organization, services, commitments, requirements, infrastructure, software, people, procedures, data).
- The control matrix exists, with every Common Criterion addressed by at least one control.
- Each control has a named owner, operating frequency, and identified evidence source.

Access and identity

- MFA is enforced for all interactive logins.
- Privileged access uses just-in-time elevation; standing administrative privileges are minimized and documented.
- Access reviews are scheduled at least quarterly with named reviewers.
- Joiner-mover-leaver process produces same-day access removal on termination.
- Conditional Access policies are documented and reviewed periodically.

Operations

- Centralized logging is in place; retention period is documented and meets regulatory and contractual obligations.
- SIEM analytics rules are in place for security events of interest; alert handling produces records suitable for evidence.
- Vulnerability management produces a continuous identification and remediation record with documented severity-based SLAs.
- Incident response runbook exists and has been exercised within the last twelve months with a documented after-action report.
- Backup and recovery testing has been performed within the observation window with measured time-to-recovery.

Change management

- Change-management policy distinguishes standard, normal, and emergency changes.
- Production deployments are tied to change records with approval, testing, and deployment evidence.
- Branch protection rules enforce code review for commits to production branches.
- Configuration baseline management and drift detection are in place for production infrastructure.

Risk and vendor management

- Annual risk assessment has been completed within the last twelve months.
- Risk register is maintained and reviewed at a documented cadence.
- Vendor inventory exists and is reconciled against accounts payable and SaaS-discovery sources.
- Material vendors have current security evidence on file (SOC 2, ISO 27001, security questionnaire).
- Subservice organization monitoring procedure is documented and operating.
- Cyber insurance coverage is in place and reviewed annually for adequacy.