

GRC Interview & Study Guide

Compliance & Audit

Prepared by **Naresh Thalla Mudhiraj** · thallanaresh.com

Personal learning project. Views are my own and do not represent any employer.

Basic to Advanced · scenario-based questions · practical examples · why each topic matters

Version 2.0 · generated 21 October 2026

Control testing, evidence, audit lifecycles, findings and continuous compliance.

Study material written in plain language. Framework and regulatory details change; always confirm against the official, current publication before relying on them.

Compliance & Audit

Control testing, evidence, audit lifecycles, findings and continuous compliance.

Q: What is the difference between compliance and security?

Basic

- Compliance means meeting specific external or internal requirements (laws, standards, contracts, policies) and being able to prove it.
- Security means actually protecting information and systems against threats.

They overlap heavily, but being compliant does not guarantee being secure: requirements are minimum baselines, may lag new threats, and audits sample evidence at a point in time. A risk-based programme uses compliance as a floor, not a ceiling.

Practical example: A company passes its PCI DSS assessment in March but stores an unencrypted export of customer data in a public cloud bucket in May. It was compliant on paper but not secure in practice.

Why this matters: A classic interview question to check you will not treat 'passing the audit' as the goal.

Q: What makes good audit evidence?

Basic

- **Relevant:** directly demonstrates the control requirement.
- **Reliable:** system-generated or from an independent source is stronger than a screenshot of a spreadsheet.
- **Complete:** covers the full population and period, not just good examples.
- **Timely and dated:** shows when the control operated.
- **Traceable:** you can follow from the population to the sample to the outcome (e.g. ticket → approval → change deployed).

Keep evidence in a central, organised repository mapped to controls.

Practical example: Good: a system-generated report of AD accounts with timestamps and the query used.
Weak: a screenshot with no date or source, or a manager's email saying 'we checked it'.

Why this matters: Weak or disorganised evidence is one of the most common reasons controls fail audits even when they actually work.

Q: What is a gap assessment?

Basic

A gap assessment compares the current state of controls against a target framework or regulation to identify what is missing or insufficient.

- Define scope and the requirements to assess against.
- Gather evidence via interviews, documents and technical review.
- Rate each requirement (e.g. met / partially met / not met).
- Prioritise gaps by risk and effort, and build a roadmap with owners and timelines.

Practical example: Before a SOC 2 project you compare current practices against the Trust Services Criteria and find 18 gaps, such as no vendor review process and no formal incident response plan. That list becomes your remediation roadmap.

Why this matters: Gap assessments are often the first deliverable in a new compliance initiative and the basis for budgeting.

Q: What is the audit lifecycle?

Basic

1. **Planning:** scope, objectives, criteria, risk assessment and audit programme.
2. **Fieldwork:** interviews, walkthroughs, evidence collection and control testing.
3. **Findings:** draft observations discussed and validated with management.
4. **Reporting:** final report with findings, ratings and management responses.
5. **Follow-up:** tracking corrective actions and validating closure.

Guidance such as ISO 19011 (management system audits) and ISACA's IT audit standards follow a similar structure.

Practical example: In an internal audit of change management, planning sets the scope (the 2 payment applications, the last 9 months), fieldwork tests a sample of emergency changes for retrospective approval, three findings are agreed with the release manager, the report is issued in week 5, and closure is validated 60 days later.

Why this matters: Knowing where you are in the lifecycle tells you what the auditor needs from you next, and helps you prepare control owners in time.

Q: What is the difference between an internal audit and an external audit?

Basic

- **Internal audit:** performed by the organisation's own audit function (or an outsourced provider acting for it), reporting to the audit committee. It covers any area the audit plan prioritises.
- **External audit:** performed by an independent third party, such as a certification body (ISO 27001), a CPA firm (SOC 2) or a QSA (PCI DSS), against a defined standard, often to give assurance to customers or regulators.
- Internal audits help prepare for external ones but do not replace them.
- Both require independence from the activities audited.

Practical example: An ISO 27001 internal audit in May finds a gap in supplier reviews, which is fixed before the external surveillance audit in September, so the certification body sees a working corrective-action process instead of a nonconformity.

Why this matters: Interviewers check that you understand who audits whom, for what purpose, and what each type of audit can and cannot give the organisation.

Q: What is a compensating control?

Basic

A compensating control is an alternative measure put in place when the primary required control cannot be implemented, which provides a similar level of protection against the same risk.

- It must address the same objective and risk as the original control.
- It should not be a control that is already required anyway.
- It must be documented, justified and tested.
- PCI DSS has a formal compensating-controls worksheet for this in the defined approach.

Practical example: A medical device cannot run anti-malware software. Compensating controls: it sits on an isolated VLAN, only the vendor's management server can reach it, application allow-listing is enabled and network traffic is monitored for anomalies.

Why this matters: Real environments always have constraints, such as legacy systems or vendor limitations. Compensating controls let you manage risk honestly rather than ignore a requirement.

Q: What is sampling in audit, and how is a sample size chosen?

Basic

Sampling means testing a subset of a population (for example 25 of 400 changes) to conclude on the whole population.

- Sample size depends on how often the control runs (annual, quarterly, monthly, weekly, daily, many times a day), the risk, and the confidence the auditor needs.
- **Audit firms typically use guidance tables:** an annual control may need 1 sample, a daily control around 20–40, and an automated control may need only 1 test plus evidence that change controls protect the configuration.
- Samples should be selected randomly or systematically across the whole period, not chosen by the auditee.
- Any exceptions found must be evaluated, and they may require a larger sample.

Practical example: For a SOC 2 Type II, the auditor asks for the complete list of 312 production changes in the period and selects 25 at random. Two lack approval evidence, so the auditor reports an exception for change management.

Why this matters: Understanding sampling helps you prepare full and accurate populations, which is one of the most common sources of audit delays and findings.

Q: A control owner says 'we do the review but never document it'. How do you explain why evidence matters and fix it?

Basic · Scenario

- **Explain simply:** for an auditor, a control without evidence did not happen. Evidence also protects the owner, because it proves they did their job.
- **Agree the minimum evidence:** who reviewed, what was reviewed, when, what was found and what action was taken.
- **Make it easy:** a ticket template, a form in the GRC tool or an automated report with an approval step.
- Set reminders in an evidence calendar so collection happens when the control runs, not at audit time.
- Check the first few cycles together.

Practical example: The firewall rule review was being done verbally in a team meeting. It is now a quarterly ticket with the rule export attached, comments on each rule changed, and the manager's approval. It takes 10 extra minutes and passes the next audit cleanly.

Why this matters: Missing evidence is one of the most common causes of audit exceptions, even when the control actually works.

Q: Explain control design effectiveness vs operating effectiveness.

Intermediate

- **Design effectiveness:** if the control is performed as described, would it actually address the risk or requirement? Tested through walkthroughs and documentation review.
- **Operating effectiveness:** did the control actually operate consistently as designed over a period of time? Tested by sampling evidence across the period.

Example: a quarterly access review is well designed if it covers all privileged accounts and requires managers to approve or revoke. It operates effectively if all four quarters were completed, by the right people, and revocations were actioned.

Practical example: Design: the change-management policy requires CAB approval for production changes (well designed). Operating: a sample of 25 changes shows 6 were deployed without approval, so the control is not operating effectively.

Why this matters: This is the core distinction behind SOC 2 Type I vs Type II and most internal control testing.

Q: What are common audit testing methods?

Intermediate

- **Inquiry:** interviewing control owners (weakest evidence on its own).
- **Observation:** watching a control being performed.
- **Inspection:** examining documents, records, configurations, tickets and logs.
- **Re-performance:** the auditor independently performs the control to check the result.

Auditors usually combine methods and select samples based on control frequency (e.g. more samples for daily controls than for annual ones).

Practical example: To test backup restoration, the auditor inquires with the infrastructure lead, inspects the backup policy and job reports, observes an engineer restoring a file server folder, and re-performs by restoring a sampled database to a test server and comparing record counts.

Why this matters: Knowing how auditors test helps you prepare the right evidence and avoid surprises.

Q: How do you prepare an organisation for an external audit or certification?

Intermediate

1. Confirm scope, framework/criteria, audit period and key dates with the auditor.
2. Perform a readiness or gap assessment against the requirements.
3. Remediate gaps and collect evidence in advance, mapped to each control.
4. Brief control owners on what to expect and how to answer (honestly, factually, within scope).
5. **Set up logistics:** single point of contact, request tracker, secure evidence sharing.
6. Track requests and turnaround daily during fieldwork.
7. Review draft findings for accuracy, agree management responses and remediation plans.

Practical example: Six months before an ISO 27001 Stage 2 audit: run a readiness assessment, close the gaps, complete one full internal audit and a management review, build an evidence library and hold mock interviews with control owners.

Why this matters: Preparation reduces findings, stress and cost — and shows leadership the GRC function adds value.

Q: How do you manage audit findings and remediation?

Intermediate

- Record each finding with its root cause, risk rating, owner, agreed action and due date.
- Focus on fixing the root cause, not just the instance found.
- Track progress regularly and escalate overdue items through governance.
- Require evidence of closure, and have it validated (by GRC or internal audit) before closing.
- **Report trends:** open/overdue findings by severity, repeat findings, ageing.

Practical example: Finding: no MFA on the admin console. Root cause: no standard for privileged access. Action plan: enforce MFA in 30 days (quick fix) and publish a privileged access standard in 90 days (systemic fix), tracked in the GRC tool and validated before closure.

Why this matters: Repeat findings damage credibility with auditors and regulators. Disciplined remediation tracking prevents them.

Q: What is control mapping / a unified control framework, and why use it?

Intermediate

Control mapping links one internal control to the requirements of several frameworks — e.g. a single access review control may satisfy ISO 27001 Annex A, SOC 2, HIPAA Security Rule, PCI DSS and NIST 800-53 requirements.

A unified (common) control framework is a single set of internal controls mapped to all applicable frameworks.

- Test once, comply many.
- Reduces audit fatigue for control owners.
- Makes it easier to add a new framework or regulation (e.g. a new national requirement) by mapping instead of starting over.

Publicly available crosswalks (e.g. NIST's online mappings) can help, but mappings should be reviewed by people who understand the context.

Practical example: One control, 'quarterly privileged access review', is mapped to ISO 27001 A.5.18, SOC 2 CC6.2/CC6.3, PCI DSS 7.2.4 and HIPAA 164.308(a)(4), so one piece of evidence serves four audits.

Why this matters: Organisations with multiple certifications cannot scale without a unified control set. This is a highly valued GRC skill.

Q: What is population completeness and accuracy, and why do auditors test it?

Intermediate

Before relying on a sample, the auditor needs confidence that the population (the full list of items, such as all leavers or all changes) is complete and accurate. This is often called IPE: information produced by the entity.

- They check the source system, the query or report used, filters and parameters, and compare totals with another source.
- A screenshot of the query and parameters, or the auditor watching the report run, is common evidence.
- If the population is incomplete, every test based on it is unreliable.

Practical example: The ticketing tool shows 140 approved production changes for the year, but the deployment pipeline logs 163 releases. The extra 23 are hotfixes pushed straight through the pipeline without a ticket. The auditor expands the population to the pipeline log and tests the hotfixes too.

Why this matters: A surprising number of audit problems come from bad populations, such as a leaver list from HR that excludes contractors.

Q: How do you write a good audit finding?

Intermediate

A good finding has five elements, often called the 5 Cs:

- **Criteria:** what should be happening, meaning the policy, standard or requirement.
- **Condition:** what was actually found, with facts and numbers.
- **Cause:** why it happened (root cause, not just a symptom).
- **Consequence:** the risk or impact to the business.
- **Corrective action:** the recommendation for what should be done.

Add a risk rating and agree the facts with management before issuing.

Practical example: Criteria: leavers' access must be removed within 1 business day. Condition: 6 of 25 sampled leavers kept access for 5–40 days. Cause: contractor leavers are not sent from the vendor-management system to IAM. Consequence: former staff could access patient data. Recommendation: automate contractor end dates into IAM.

Why this matters: Clear, evidence-based findings are accepted and fixed faster. Vague findings cause arguments and delays.

Q: What is a root cause analysis, and which techniques are used?

Intermediate

Root cause analysis (RCA) finds the underlying reason a problem happened so it can be fixed permanently, rather than fixing only the symptom.

- **5 Whys:** keep asking 'why' until you reach a process or design cause.
- **Fishbone (Ishikawa) diagram:** group possible causes into people, process, technology, environment, and so on.
- **Timeline analysis:** for incidents and repeated failures.
- A good root cause is something you can act on, such as a missing process, unclear ownership or a design flaw. 'Human error' alone is rarely the real root cause.

Practical example: Why were patches late? The team missed the window. Why? There was no maintenance window agreed for clinical servers. Why? Clinical leadership was never asked. The root cause is a missing governance agreement, not lazy admins.

Why this matters: Without RCA, findings come back in the next audit. Regulators and certification bodies look for repeat findings as a sign of weak management.

Q: An auditor asks for evidence you do not have. What do you do?

Intermediate · Scenario

- Be honest. Never create or back-date evidence. Fabricated evidence is far worse than a gap.
- Clarify exactly what the auditor needs. Sometimes alternative evidence exists, such as system logs, tickets or meeting minutes.
- If the control did not operate, say so, explain any compensating controls and share the remediation plan.
- Agree a timeline for anything that can legitimately be provided later.
- Afterwards, fix the evidence process so the gap does not happen again (evidence calendar, automated capture).

Practical example: The auditor asks for sign-off of the Q2 access review. There is no signed form, but the ticketing system shows the reviewer's completed task with a timestamp and the list of removals. The auditor accepts it, and a signed workflow is added for future reviews.

Why this matters: Auditors judge your integrity as well as your controls. How you handle a gap shapes how much they trust everything else you show them.

Q: You have three audits (ISO 27001 surveillance, SOC 2 Type II, and a customer audit) in the same quarter. How do you manage it?

Intermediate · Scenario

- Build one combined audit calendar and avoid overlapping fieldwork weeks for the same teams.
- **Use your unified control framework:** collect evidence once and reuse it across audits.
- Create a single evidence repository with clear naming and dates.

- **Prepare control owners:** brief them on what each auditor will ask and when.
- For the customer audit, offer existing reports (SOC 2, ISO certificate) first to reduce scope.
- Hold a short daily stand-up during fieldwork to clear requests quickly.

Practical example: By sharing 70% of the evidence across all three audits and reducing the customer audit to a review of the SOC 2 report plus a 2-hour call, the team completes the quarter without pulling engineers off product work for weeks.

Why this matters: Audit fatigue is a real risk to operations and staff morale. Good coordination protects the business and the audit outcomes.

Q: The same audit finding has been raised three years in a row. What do you do differently this time?

Intermediate · Scenario

- Treat repeat findings as a governance issue, not just a technical one. Report it to the risk or audit committee with its history.
- Do a proper root cause analysis. Previous fixes probably addressed symptoms (one-off clean-ups) rather than the process.
- Assign a senior owner with authority and budget, and break the fix into milestones with dates.
- Validate closure with testing over a period, not just a single point in time.
- Add a KRI so drift is visible before the next audit.

Practical example: Excess admin rights had been found 3 years running because each year IT cleaned up just before the audit. This time a privileged access management tool with just-in-time access is funded, and a monthly KRI tracks standing admin accounts.

Why this matters: Repeat findings damage credibility with auditors and regulators and suggest weak management oversight.

Q: What is continuous control monitoring (CCM)?

Advanced

CCM uses automation to test controls continuously or frequently instead of only during annual audits. Examples:

- Automated checks of cloud configuration against benchmarks.
- Daily comparison of HR leavers with active accounts.
- Alerts when endpoints are missing encryption or EDR.
- GRC platform integrations that pull evidence directly from source systems.

Benefits: earlier detection of control failures, less manual evidence collection, and a more accurate real-time compliance posture.

Practical example: Instead of sampling quarterly, a script checks all cloud storage buckets every hour for public access and opens a ticket automatically, so evidence for the auditor is a full population with timestamps.

Why this matters: Moving from point-in-time audits to continuous assurance is a major trend and a strong talking point for senior roles.

Q: What is the difference between an attestation (for example SOC 2) and a certification (for example ISO 27001)?

Advanced

- **Certification:** an accredited certification body confirms that a management system conforms to a standard and issues a certificate, usually valid for 3 years with surveillance audits.
- **Attestation:** an independent CPA firm issues an opinion on management's description and controls against criteria, such as the Trust Services Criteria for SOC 2. The output is a detailed report, not a certificate.
- A SOC 2 report shows tests and results, including exceptions. An ISO certificate shows that the scope is certified without control-level test details.
- **Customers often want both:** the certificate for the management system and the SOC 2 report for control detail.

Practical example: A US customer asks a SaaS vendor for its 'SOC 2 certificate'. The GRC team explains that SOC 2 is an attestation report, shares it under NDA, and adds the ISO 27001 certificate for the management system.

Why this matters: Mixing up the two is a common interview mistake. Knowing the difference helps you choose the right assurance for customers and interpret vendors' evidence correctly.

Q: Leadership has bought a GRC platform and asks you to roll it out. How do you make it succeed?

Advanced · Scenario

- Fix the process before you automate it. A tool will not fix unclear ownership or a messy control library.
- Start with a unified control set and clear data model (risks, controls, assets, vendors, evidence).
- **Prioritise high-volume, repetitive work:** evidence collection through integrations (cloud, identity, ticketing), questionnaires and reminders.
- **Phase the rollout:** one framework or module first, measure the time saved, then expand.
- **Plan adoption:** training, dashboards people actually use, and executive sponsorship.
- Keep exports and data ownership in mind to avoid vendor lock-in.

Practical example: A company moves from spreadsheets to a GRC platform, starting with ISO 27001 and SOC 2 evidence from its cloud and identity systems. After 6 months, evidence-gathering time for the SOC 2 audit drops by roughly half, which justifies adding the vendor-risk module next.

Why this matters: Many GRC tool projects fail because they automate chaos. Interviewers value people who can deliver a working rollout.

Q: A regulator sends an information request with a 10-day deadline. How do you respond?

Advanced · Scenario

1. Log it, confirm the exact deadline and inform legal and compliance leadership immediately.
2. Clarify the scope. If anything is ambiguous, ask the regulator early and in writing.
3. Assign owners for each item and a single coordinator. All communication goes through the agreed channel.
4. Gather evidence, review it for accuracy and consistency, and have legal review the response before it is sent.
5. If the deadline cannot be met for some items, request an extension early with a reason.
6. Keep a complete record of what was sent and when, and track any follow-up commitments.

Practical example: A health regulator asks for the last 12 months of access-review evidence for the EHR and the incident log. The GRC lead coordinates responses from IT and privacy, legal reviews them, and the pack is submitted on day 8 with a cover letter that indexes every document.

Why this matters: Regulatory responses are high-stakes. Accuracy, consistency and timeliness matter as much as the content.

Q: You discover a control has silently failed for six months (for example a log-forwarding agent stopped). What do you do?

Advanced · Scenario

- Fix it immediately and confirm the control works again.
- **Assess the impact of the gap:** what activity went unmonitored, and could an incident have been missed? Review other sources for that period if possible.
- Record it as an issue or deficiency, and tell the right people: the risk owner and, if it affects an audit period, the auditor.
- Do a root cause analysis. Why did nothing alert you? Add monitoring of the control itself (for example, an alert when log volume drops).
- Update the risk assessment and consider whether customers or regulators must be informed, based on contracts and law.

Practical example: The SIEM stopped receiving logs from 30 servers after an agent update. After the fix, a heartbeat alert is added so any source silent for more than 4 hours opens a ticket, and the SOC 2 auditor is told of the gap and its remediation.

Why this matters: Controls that fail silently are worse than controls that fail loudly. Monitoring the health of controls is a mark of a mature programme.