GRC Interview & Study Guide

GRC Fundamentals

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

What GRC is, how the pieces fit together and the core vocabulary every GRC professional uses.

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

GRC Fundamentals

What GRC is, how the pieces fit together and the core vocabulary every GRC professional uses.

Q: What is GRC and why do organisations need it?

Basic

GRC stands for Governance, Risk and Compliance. It is an integrated way of making sure an organisation is directed properly, understands the risks to its objectives and meets its legal, regulatory and contractual obligations.

- **Governance:** who decides, how decisions are made, and how leadership sets direction and holds people accountable.
- **Risk management:** identifying, assessing, treating and monitoring things that could stop the organisation achieving its objectives.
- **Compliance:** demonstrating that the organisation follows the laws, regulations, standards, contracts and internal policies that apply to it.

Doing the three together avoids duplicated effort (one control can satisfy many requirements), gives leadership one view of risk, and makes security and compliance part of normal business decisions rather than a last-minute exercise.

Practical example: A hospital group runs separate ISO 27001, HIPAA and payment-card projects with three teams collecting the same access-review evidence. Moving to one GRC programme with a shared control set and a single risk register cuts the evidence requests by about two-thirds and gives the board one cyber risk report instead of three.

Why this matters: Interviewers use this to check you see GRC as a business enabler, not a paperwork function. A joined-up GRC programme reduces cost, audit fatigue and surprises.

Q: What is the difference between a threat, a vulnerability and a risk?

Basic

- **Threat:** something that could cause harm — e.g. a ransomware group, an insider, a flood.
- **Vulnerability:** a weakness that a threat could exploit — e.g. an unpatched server, weak passwords, no offsite backup.
- **Risk:** the combination of the likelihood that a threat exploits a vulnerability and the impact if it does — e.g. 'Ransomware encrypts the unpatched billing server, stopping invoicing for several days.'

A useful way to write a risk statement is: cause (threat + vulnerability) → event → consequence (impact on objectives).

Practical example: Threat: a ransomware gang. Vulnerability: an internet-facing VPN appliance missing a critical patch. Risk: the chance that the gang exploits the VPN and encrypts the patient scheduling system, and the impact on clinics if it does.

Why this matters: Clear risk statements are the foundation of a useful risk register. Vague entries like 'cyber risk' cannot be owned, measured or treated.

Q: What is the CIA triad?

Basic

The CIA triad is the core model for information security objectives:

- **Confidentiality:** only authorised people can access information (encryption, access control).
- **Integrity:** information is accurate and has not been altered without authorisation (hashing, change control, segregation of duties).

- **Availability:** information and systems are accessible when needed (redundancy, backups, DR).

Many organisations also consider authenticity, non-repudiation and, for data protection, privacy.

Practical example: Patient records: confidentiality is breached if a receptionist views a celebrity's file; integrity is breached if a medication dose is altered; availability is breached if the EHR is down during a night shift.

Why this matters: Impact assessments, data classification and control selection all come back to which of these properties matters most for a given asset.

Q: What is a control, and what types of controls are there?

Basic

A control is any measure that modifies risk — a policy, process, device, practice or other action.

By function:

- **Preventive:** stop an event happening (MFA, firewalls, segregation of duties).
- **Detective:** identify an event that has happened (SIEM alerts, log reviews, reconciliations).
- **Corrective:** limit damage and restore (backups, incident response, patching after an event).
- **Others often mentioned:** deterrent, compensating and directive controls.

By nature:

- Administrative / managerial (policies, training), technical / logical (encryption, access control) and physical (locks, CCTV).

Practical example: For laptop theft: full-disk encryption (preventive, technical), asset register reconciliation (detective, administrative), remote wipe (corrective, technical) and a locked storage cabinet (preventive, physical).

Why this matters: A balanced control set uses several types together (defence in depth). Auditors and assessors will ask how each risk is prevented, detected and corrected.

Q: What does 'defence in depth' mean?

Basic

Defence in depth means layering several independent controls so that if one fails, others still protect the asset.

- **Example layers:** security awareness → email filtering → endpoint protection → least privilege → network segmentation → monitoring → backups.

It avoids relying on a single point of failure and combines preventive, detective and corrective controls.

Practical example: A phishing email gets past the mail filter (layer 1), but the user's MFA blocks the stolen password (layer 2), EDR flags the payload (layer 3) and network segmentation stops lateral movement (layer 4).

Why this matters: Breaches usually happen when several controls fail together. Layering controls makes that much less likely.

Q: What is an information asset, and why does GRC start with an asset inventory?

Basic

An information asset is anything that has value to the organisation and handles information: data sets, applications, servers, cloud services, laptops, people with key knowledge and even paper records.

- You cannot protect, risk-assess or audit what you do not know exists.
- The inventory records the owner, classification, location, business process supported and criticality of each asset.
- It feeds the risk assessment, BIA, access reviews, vulnerability scanning scope and incident response.

ISO 27001:2022 control 5.9 (Inventory of information and other associated assets) and NIST CSF 2.0 ID.AM both expect it.

Practical example: A discovery scan finds 40 cloud VMs that are not in the CMDB, including a staging database with a copy of live patient data. Adding them to the inventory with owners brings them into patching, backup and access-review scope.

Why this matters: Most audit findings and breaches trace back to an unknown or unowned asset, such as a forgotten test server or an unmanaged SaaS app. The inventory is the foundation every other control builds on.

Q: What is data classification, and what levels are commonly used?

Basic

Data classification labels information by how sensitive it is so that people apply the right level of protection.

- **A common four-level scheme:** Public, Internal, Confidential and Restricted (or Highly Confidential).
- Each level has handling rules for storage, sharing, encryption, printing, retention and disposal.
- The data owner decides the level; users apply labels; tools (DLP, sensitivity labels) enforce them.
- Keep it simple. If staff need a flowchart to choose a label, the scheme is too complex.

Practical example: Marketing brochures are Public, the internal org chart is Internal, customer contracts are Confidential and patient lab results are Restricted, so they must be encrypted, never emailed externally without approval, and kept in the approved EHR only.

Why this matters: Without classification, everything is protected the same way. That means too much cost for public data and too little protection for patient records or card data.

Q: What is the principle of least privilege, and how is it different from need-to-know?

Basic

Least privilege means giving users, services and systems only the minimum access rights they need to do their job, and only for as long as they need them.

Need-to-know is narrower: even if someone has a clearance or a role, they only see specific information when their task requires it.

- Least privilege is about permissions (what you can do).
- Need-to-know is about information (what you can see).
- Both reduce the damage an attacker or careless insider can do with one compromised account.

Practical example: A billing clerk needs to read invoices but not change the payment bank details. Removing their edit right on vendor master data (least privilege) closes a common payment-fraud path. A nurse can access the EHR but only sees records of patients on her ward (need-to-know).

Why this matters: Excessive access is one of the most frequent audit findings and one of the main ways attackers move through a network after the first compromise.

Q: What is segregation of duties (SoD) and why does it matter?

Basic

Segregation of duties splits a sensitive task so that no single person can complete it end to end and hide an error or fraud.

- **Classic split:** request, approve, execute and reconcile should sit with different people.
- **In IT:** the developer who writes code should not be the one who approves and deploys it to production.
- Where staff numbers are too small for full SoD, use compensating controls: management review of logs, rotation, or automated alerts.

SoD conflicts are usually managed with a conflict matrix in the ERP or IAM tool.

Practical example: In a small clinic the same IT admin creates user accounts and reviews the access list. As a compensating control, the practice manager reviews a monthly system-generated report of new accounts and signs it off.

Why this matters: SoD is a key fraud-prevention control that external and SOX-style auditors test first. A single person with end-to-end power is both a fraud risk and an error risk.

Q: What is the difference between authentication, authorisation and accounting (AAA)?

Basic

- **Authentication:** proving who you are, such as a password plus a one-time code, a smart card or a passkey.
- **Authorisation:** deciding what you are allowed to do once authenticated (roles, permissions, policies).
- **Accounting (or auditing):** recording what you actually did, such as logins, data accessed and changes made, so actions can be traced back to a person.

Together they give you identity assurance, controlled access and accountability. Weakness in any one of them creates gaps auditors will find.

Practical example: A doctor logs in with SSO and MFA (authentication). The EHR role lets her prescribe but not change billing (authorisation). Every record she opens is logged, so a privacy officer can investigate a complaint later (accounting).

Why this matters: AAA is the vocabulary of every access-control conversation and maps directly to ISO 27001 controls 5.15–5.18, 8.5 and 8.15.

Q: A colleague in finance receives an urgent email from the 'CEO' asking for a same-day bank transfer to a new supplier. Which GRC concepts and controls apply?

Basic · **Scenario**

This is business email compromise (BEC), a social-engineering threat that exploits urgency and authority.

- **Preventive controls:** a payment policy requiring call-back verification on a known number for new or changed bank details, dual approval for payments above a threshold, and email authentication (SPF, DKIM, DMARC).
- **Detective controls:** banner warnings on external email, monitoring for look-alike domains, finance reconciliation.
- **People controls:** targeted awareness for finance and executive assistants.
- **Response:** if money has moved, contact the bank immediately to attempt a recall, and treat it as a security incident.

It shows why technical, administrative and people controls must work together (defence in depth).

Practical example: The finance officer follows the call-back rule, phones the CEO's known mobile number and finds the request is fake. The incident is reported, the look-alike domain is blocked and the case becomes a short awareness story for all staff.

Why this matters: BEC losses are often larger than ransomware losses for many organisations, and they are prevented mostly by simple process controls rather than expensive technology.

Q: What is the difference between a policy, standard, procedure and guideline?

Intermediate

- **Policy:** high-level statement of management intent and direction (what and why). Approved by senior management, changes rarely. Mandatory.
- **Standard:** specific mandatory requirements that support a policy — e.g. 'passwords must be at least N characters; MFA required for remote access'.
- **Procedure:** step-by-step instructions for performing a task (how) — e.g. the user access provisioning procedure.
- **Guideline:** recommended practice or advice. Not mandatory, gives flexibility.

Together they form a document hierarchy: policies set direction, standards make it measurable, procedures make it repeatable.

Practical example: Policy: 'All remote access must use MFA.' Standard: 'MFA must use an authenticator app or FIDO2 key; SMS is not allowed.' Procedure: the service-desk steps to enrol a user. Guideline: tips on choosing a hardware key.

Why this matters: Auditors test against this hierarchy. A policy without standards and procedures is hard to implement consistently and hard to evidence.

Q: Explain the Three Lines Model.

Intermediate

The Three Lines Model (updated by the Institute of Internal Auditors in 2020, previously 'Three Lines of Defence') describes roles in governance and risk management:

- **Governing body:** accountable to stakeholders for oversight.
- **First line:** operational management that owns and manages risk day to day (e.g. IT operations, business units).
- **Second line:** specialist functions that provide expertise, frameworks, monitoring and challenge — e.g. risk management, compliance, information security GRC.
- **Third line:** internal audit, providing independent and objective assurance to the governing body.

External auditors and regulators sit outside the model as additional assurance providers.

Practical example: First line: the IT operations team that runs access reviews. Second line: the GRC/risk team that sets the access standard and checks the reviews happen. Third line: internal audit, which independently tests the whole thing and reports to the audit committee.

Why this matters: Most GRC analyst roles sit in the second line. Knowing where you sit explains why you advise and challenge rather than own operational controls.

Q: What is the difference between inherent risk and residual risk?

Intermediate

- **Inherent risk:** the level of risk before taking existing controls into account (or assuming no controls).
- **Residual risk:** the level of risk that remains after controls are applied.
- **Some organisations also track target risk:** the level they aim to reach after planned treatments.

Example: storing patient data on internet-facing systems has a high inherent risk. With MFA, encryption, monitoring and patching in place, the residual risk may be medium. If that is still above appetite, more treatment is planned.

Practical example: An unpatched public web server might be rated High inherently. After adding a WAF, weekly patching and monitoring, the residual risk is rated Medium, and the owner formally accepts that residual level.

Why this matters: Comparing inherent and residual risk shows how much value your controls deliver and where to focus control testing.

Q: What is the difference between a framework, a standard, a regulation and a contract requirement?

Intermediate

- **Framework:** a structured set of good practices you can adopt and adapt, such as NIST CSF, COBIT or the NIST AI RMF. Usually voluntary.
- **Standard:** a formal, published specification you can be audited or certified against, such as ISO/IEC 27001 or PCI DSS.
- **Regulation or law:** mandatory rules made by a government or regulator, such as GDPR, HIPAA, Singapore's PDPA or India's DPDP Act. Breaking them brings penalties.
- **Contract requirement:** obligations you agreed to with a customer or partner, such as 'maintain SOC 2 Type II' or 'notify us of breaches within 24 hours'.

A good obligations register lists all four, because contracts are often stricter than the law.

Practical example: A health-tech company follows NIST CSF voluntarily, certifies to ISO 27001, must obey national health-data law and its health regulator's requirements, and has a hospital contract demanding 24-hour breach notice, which is tighter than any law it falls under.

Why this matters: Knowing the source of each requirement tells you how much flexibility you have and what the consequence of failure is: certification loss, a fine or a breach of contract.

Q: What is an obligations register (compliance universe) and how do you build one?

Intermediate

An obligations register is a single list of every legal, regulatory, contractual and voluntary requirement that applies to the organisation, with an owner and the controls that meet it.

1. **Identify sources:** laws in each country of operation, sector regulators, standards you certify to, key customer contracts and internal policies.
2. Break each source into specific, testable obligations.
3. Assign an owner and map each obligation to controls in your control library.
4. Record evidence, testing frequency and status.
5. Set up horizon scanning so new or changed obligations are added (regulatory change management).

Practical example: When expanding into a new country, a company adds the local data protection law, the national health-data law, the health regulator's requirements and two hospital contracts to its register, then maps them to existing ISO 27001 controls and finds 6 new obligations, mostly data localisation and local notification.

Why this matters: Regulators and auditors increasingly ask 'how do you know what applies to you?'. The register is the answer, and it stops obligations being missed when people leave.

Q: What is security by design and privacy by design, and how does GRC support them?

Intermediate

Security and privacy by design mean building protections into systems, products and processes from the start, rather than bolting them on after go-live.

- **GRC embeds checkpoints into the project lifecycle:** a security and privacy screening at idea stage, a risk assessment or DPIA at design, security testing before release.
- Requirements come from the obligations register and the data classification.

- Defaults should be the most protective setting (privacy by default, GDPR Article 25).

ISO 27001:2022 controls 5.8 (information security in project management) and 8.25–8.29 (secure development) support this.

Practical example: A product team plans a patient app that uploads location data. An early privacy screening triggers a DPIA, and the design changes to collect location only when the user requests a nearby-clinic search, with data deleted after 24 hours.

Why this matters: Fixing a design flaw after launch can cost many times more than addressing it at design stage, and regulators expect proof that privacy was considered up front.

Q: What is the difference between a control objective, a control activity and a control test?

Intermediate

- **Control objective:** what you want to achieve. Example: 'Only authorised users have access to production systems.'
- **Control activity:** the specific action that achieves it. Example: 'The application owner reviews all production access quarterly and removes access that is no longer needed.'
- **Control test:** how you check it works. Example: 'Select 2 of 4 quarterly reviews, confirm they were completed on time, and reperform removal for 10 flagged users.'

Writing controls at this level of precision is what makes them auditable.

Practical example: An auditor rejects the control 'backups are performed'. It is rewritten as 'Daily automated backups of the EHR database are monitored; failures generate a ticket resolved within 24 hours; a restore test is performed quarterly', which can be tested.

Why this matters: Vague controls ('access is managed appropriately') cannot be tested, so they lead to findings. Precise wording saves time for owners and auditors alike.

Q: You discover a team using an unapproved cloud file-sharing app (shadow IT) with customer data. What do you do?

Intermediate · **Scenario**

1. **Assess the immediate risk:** what data, how much, who has access, are links shared publicly?
2. **Contain sensibly:** restrict public links and stop new uploads of sensitive data, without shutting down the team's work overnight.
3. Understand why they used it. Usually the approved tool is missing a feature or is too slow.
4. **Decide:** bring the app under governance (vendor assessment, SSO, contract) or migrate the data to an approved tool.
5. **Fix the root cause:** a quick request process for new tools, plus CASB or SaaS discovery to spot it earlier.

Treat it as a learning moment, not a punishment, unless policy was knowingly broken.

Practical example: A sales team shares proposals through a free file-sharing account because the approved tool blocks large files. GRC removes public links, the IT team raises the size limit in the approved tool, and the data is migrated within two weeks. A SaaS discovery tool is added to catch the next case early.

Why this matters: Shadow IT is a symptom of unmet business needs. A response that only punishes drives it further underground.

Q: A manager asks you to mark a control as 'effective' because the audit is next week, although you know it is not working. What do you do?

Intermediate · **Scenario**

- Decline politely and clearly. Misstating control status is a matter of integrity and can expose you and the organisation to serious consequences.
- **Explain the risk:** auditors usually find it anyway, and a concealed failure turns a minor finding into a credibility problem.
- **Offer a constructive path:** record the control as 'not effective' or 'partially effective', document the remediation plan with dates, and show any compensating control.
- If pressure continues, escalate through the proper channel (CISO, compliance head, ethics line).

ISACA's Code of Professional Ethics and most certification bodies expect exactly this behaviour.

Practical example: The quarterly access review for a finance system was skipped. You record it as a deficiency, complete a catch-up review before the audit, and show the auditor both the gap and the fix. The auditor raises only a minor observation instead of a finding about concealment.

Why this matters: Your value as a GRC professional depends on people trusting your reports. Being honest about a gap with a plan is always better than hiding it.

Q: You are asked to classify a new data set that mixes patient names, appointment dates and anonymous survey scores. How do you decide the level?

Intermediate · **Scenario**

- Classify by the most sensitive element. Patient names linked to appointment information are health-related personal data, so the whole set takes the highest relevant level (for example, Restricted).
- Check whether the sensitive fields are really needed. If analysts only need the survey scores, provide a de-identified extract instead.
- Confirm with the data owner, who makes the final decision.
- **Apply the handling rules:** approved storage, access only for named roles, encryption, retention limits.
- Note that 'anonymous' survey data can become identifiable once it is combined with names and dates.

Practical example: Analysts originally asked for the full extract. After the review they receive only survey scores with clinic and month. That is classified Internal, while the full data set stays Restricted in the EHR reporting zone.

Why this matters: Aggregated data sets are often under-classified because each field looks harmless on its own. The combination is what creates the risk.

Q: What is zero trust, and how would you explain its impact on GRC controls?

Advanced

Zero trust is a security model that assumes no user, device or network is trusted by default. Every access request is verified explicitly, using identity, device health and context, and access is limited and continuously re-evaluated.

- **Core ideas (NIST SP 800-207):** verify explicitly, use least-privilege access, assume breach.
- It shifts controls from the network perimeter (VPN, firewall zones) to identity, device and data.
- **GRC impact:** control descriptions, risk assessments and audit tests change. For example, 'is the device compliant before access?' replaces 'is the user on the corporate network?'.
- It is a journey, not a product. Maturity models such as CISA's Zero Trust Maturity Model help plan and measure progress.

Practical example: A hospital retires its flat VPN. Clinicians now reach the EHR through an identity-aware proxy that checks MFA, device encryption and EDR status on every session. The access-control policy and the SoA justifications are updated to describe the new controls.

Why this matters: Many organisations claim 'zero trust' without changing their controls. GRC professionals need to test whether the principles are actually implemented and reflected in policy.

Q: What is the difference between enterprise risk management (ERM) and IT/cyber risk management, and how should they connect?

Advanced

- **ERM covers all risks to business objectives:** strategic, financial, operational, legal, reputational. It is usually owned by a CRO or risk committee and uses frameworks such as COSO ERM or ISO 31000.
- Cyber risk management covers risks from information and technology, using frameworks such as ISO 27005, NIST SP 800-30 or FAIR.
- They should connect through a common risk taxonomy, shared impact scales and escalation rules, so top cyber risks appear in the enterprise register in business language.
- Avoid two disconnected registers with different scoring. The board then cannot compare a cyber risk with a market or credit risk.

Practical example: A cyber risk 'ransomware disrupts clinical systems' is linked in the ERM register to the enterprise risk 'inability to deliver patient care', using the same 5-point impact scale for patient safety, finance and reputation.

Why this matters: Boards decide on enterprise risk, so cyber risk only gets proper attention and funding when it is expressed in the same terms as other business risks.

Q: A new CEO says 'compliance is just paperwork, cut the GRC budget by half'. How do you respond?

Advanced · **Scenario**

- Do not argue in compliance jargon. Translate GRC into what the CEO cares about: revenue, licence to operate, customer trust and avoiding surprises.
- **Show the value:** deals won because of ISO 27001 or SOC 2, regulatory exposure avoided, audit findings closed, and the cost of questionnaires handled faster.
- **Show the efficiencies already made, and offer more:** automating evidence, rationalising duplicate frameworks into one control set, retiring low-value activities.
- **Present options with consequences:** 'With 50% we can keep ISO 27001 and HIPAA but must drop SOC 2, which three enterprise customers require.'
- Let leadership make an informed decision, and record it as a risk acceptance if they reduce coverage.

Practical example: You show the CEO that 4 of the last 6 enterprise contracts required SOC 2 evidence, worth a large share of new revenue, and propose saving 20% through a GRC tool that automates evidence, rather than cutting the certification that wins deals.

Why this matters: GRC teams survive budget pressure by proving business value and offering choices, not by defending activity for its own sake.