

GRC Interview & Study Guide

Business Continuity & Incident Management

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

Keeping the business running and responding well when things go wrong: BIA, RTO/RPO, DR, incident response and lessons learned.

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

Business Continuity & Incident Management

Keeping the business running and responding well when things go wrong: BIA, RTO/RPO, DR, incident response and lessons learned.

Q: What is the difference between business continuity, disaster recovery and incident response?

Basic

- **Business continuity (BC):** keeping critical business functions running during and after a disruption — people, processes, facilities, suppliers and technology. ISO 22301 is the main standard.
- **Disaster recovery (DR):** restoring IT systems, infrastructure and data after a disruption. DR is a subset of BC focused on technology.
- **Incident response (IR):** detecting, containing, eradicating and recovering from security incidents (e.g. malware, data breach).

They connect: a major cyber incident may trigger IR, then DR to restore systems, under the overall BC and crisis management plan.

Practical example: Ransomware hits: incident response contains and eradicates it; DR restores servers from immutable backups; business continuity has clinics switch to paper downtime procedures until the EHR is back.

Why this matters: Interviewers check you understand how these plans fit together rather than treating them as separate documents.

Q: What is a Business Impact Analysis (BIA)?

Basic

A BIA identifies critical business processes and the impact over time of their disruption. It typically determines:

- Impacts (financial, operational, regulatory, reputational, patient/customer safety) at different time intervals.
- Maximum Tolerable Period of Disruption (MTPD, also called maximum tolerable downtime).
- Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO).
- **Dependencies:** systems, people, suppliers, facilities and data.
- Minimum resources needed to resume.

The BIA drives continuity strategy and DR investment.

Practical example: The BIA shows outpatient scheduling can be down for 24 hours before revenue and patient-safety impacts become severe, but medication dispensing can only tolerate 2 hours. That sets very different recovery priorities.

Why this matters: Without a BIA, recovery priorities are guesswork and DR spending may protect the wrong systems.

Q: Explain RTO, RPO and MTPD.

Basic

- **Recovery Time Objective (RTO):** the target time to restore a process or system after disruption.
- **Recovery Point Objective (RPO):** the maximum acceptable amount of data loss, measured in time — e.g. an RPO of 1 hour means backups or replication must allow restoring to no more than 1 hour before the incident.
- **Maximum Tolerable Period of Disruption (MTPD):** the point after which the impact becomes unacceptable. RTO must be shorter than MTPD.

Shorter RTO and RPO usually mean higher cost (e.g. real-time replication vs nightly backups), so they should be justified by the BIA.

Practical example: Pharmacy system: RTO 2 hours (back up within 2 hours), RPO 15 minutes (lose at most 15 minutes of data), MTPD 8 hours (beyond that, patient harm becomes unacceptable).

Why this matters: These metrics translate business tolerance into technical requirements — and are a guaranteed interview question.

Q: What are the phases of incident response?

Basic

NIST SP 800-61 has long described the incident response lifecycle as:

1. **Preparation:** policies, plans, roles, tools, training, contact lists.
2. **Detection and analysis:** identify, validate, classify and prioritise incidents.
3. **Containment, eradication and recovery:** limit damage, remove the cause, restore normal operations.
4. **Post-incident activity:** lessons learned and improvements.

SP 800-61 Revision 3 (2024) re-frames incident response around the NIST CSF 2.0 functions, but the lifecycle above is still widely used. Other models (e.g. SANS PICERL) are similar.

Practical example: Phishing-based account takeover: prepare (playbook, logging), detect (impossible-travel alert), contain (reset credentials, revoke sessions), eradicate (remove mail rules), recover (restore mailbox), lessons learned (enforce phishing-resistant MFA).

Why this matters: A structured lifecycle prevents chaotic responses and supports regulatory reporting and evidence handling.

Q: What is a business continuity plan (BCP), and what should it contain?

Basic

A BCP documents how the organisation will continue critical activities during and after a disruption. Typical contents:

- Scope, objectives and the critical activities covered (from the BIA).
- Roles and responsibilities, including the crisis or incident management team and deputies.
- Activation criteria and escalation steps.
- Recovery strategies and workarounds for each critical activity (alternate site, manual processes, alternate suppliers).
- Contact lists, communication plans and resources needed.
- Testing, maintenance and review arrangements.

ISO 22301 sets requirements for a business continuity management system.

Practical example: A clinic's BCP for an EHR outage includes paper downtime forms, a printed list of the day's appointments generated every evening, a phone tree and the rule that IT must confirm data reconciliation before normal operations resume.

Why this matters: A plan that is short, current and practised is far more useful in a crisis than a large document nobody has read.

Q: What is a tabletop exercise?

Basic

A tabletop exercise is a discussion-based session where the team walks through a realistic scenario and talks about how they would respond, without touching live systems.

- A facilitator presents the scenario in stages (injects), such as 'the attackers now threaten to publish data'.
- Participants explain decisions, roles and communications.
- It reveals gaps in plans, contacts, decision rights and coordination.
- It ends with an after-action report and improvement actions.

Practical example: A 2-hour ransomware tabletop with executives finds that nobody knows who can approve shutting down the patient portal, and that the legal team does not have the cyber insurer's hotline. Both are fixed within a week.

Why this matters: Tabletops are a low-cost, high-value way to test plans and train leaders, and regulators and insurers increasingly ask for evidence of them.

Q: What is the 3-2-1 backup rule, and why are immutable backups important?

Basic

- **3-2-1:** keep at least 3 copies of data, on 2 different media or storage types, with 1 copy off-site.
- **Many now extend it to 3-2-1-1-0:** one copy offline or immutable, and zero errors after restore verification.
- Immutable backups cannot be changed or deleted for a set period, even by administrators, which protects them against ransomware and malicious insiders.
- Backups must be tested by actual restores. An untested backup is only a hope.

Practical example: After ransomware encrypts file servers, the attackers also delete the online backups. The organisation restores from an immutable cloud copy locked for 30 days and is back online in 2 days without paying.

Why this matters: Attackers deliberately target backups first. Immutable, tested backups are often what decides whether an organisation recovers quickly or pays a ransom.

Q: A user reports a suspicious email after clicking the link. What should happen next?

Basic · Scenario

1. Thank the user. Fast reporting is what you want, so never punish someone who reports.
2. The security team checks whether credentials were entered or a file was run.
3. **If credentials were entered:** reset the password, revoke sessions, check MFA prompts and mailbox rules, and review sign-in logs.
4. Search for the same email in other mailboxes and remove it.
5. Block the sender and malicious URLs.
6. Record the incident and use it (anonymised) in awareness.

Practical example: The user reported within 3 minutes. The SOC finds the same email in 40 mailboxes, removes it from all of them, resets the user's password and blocks the domain, all within 20 minutes.

Why this matters: Speed of reporting decides how much damage phishing does. A reporting culture is built by how the first report is handled.

Q: How should business continuity and DR plans be tested?

Intermediate

Common test types, from least to most disruptive:

- **Plan review / checklist:** owners review the plan for accuracy.
- **Tabletop exercise:** discussion-based walkthrough of a scenario with key decision makers.
- **Simulation / functional exercise:** practising specific procedures under realistic conditions.
- **Parallel test:** recovery systems are brought up alongside production.
- **Full interruption test:** production is actually failed over (highest assurance, highest risk).

After each test, document results, gaps and actions, and update plans. Test at least annually and after significant changes.

Practical example: Year plan: Q1 tabletop for executives on ransomware, Q2 DR failover test of the EHR to the secondary site, Q3 call-tree test, Q4 full simulation. Each one produces an after-action report with tracked improvements.

Why this matters: An untested plan is an assumption. Testing proves recovery objectives can be met and builds team confidence.

Q: What is the difference between a security event, an incident and a breach?

Intermediate

- **Event:** any observable occurrence in a system or network (e.g. a failed login).
- **Incident:** an event or series of events that actually or potentially jeopardises confidentiality, integrity or availability, or violates security policy.
- **Breach (data breach):** an incident that results in unauthorised access to, disclosure, or loss of protected data, typically triggering legal notification analysis (e.g. HIPAA, GDPR, DPDP).

Not every incident is a breach, and the legal definition of 'breach' varies by law, so legal and privacy teams should be involved in the decision.

Practical example: Event: a failed login. Incident: 500 failed logins followed by a successful one from an unusual country. Breach: confirmed access to a mailbox containing patient information, which triggers notification analysis.

Why this matters: Getting the classification right determines escalation, notification deadlines and potential regulatory exposure.

Q: What is the role of GRC during a major security incident?

Intermediate

GRC is usually not the hands-on technical responder, but plays important roles:

- Making sure the incident response plan, roles and escalation paths are followed.
- Supporting the breach/notification assessment with legal and privacy (regulatory timelines, contractual notification duties to customers).
- Tracking decisions and evidence for regulators, auditors and insurers.
- Coordinating third-party notifications and obligations.
- **After the incident:** running lessons learned, updating the risk register, tracking corrective actions and reporting to leadership.

Practical example: During a ransomware incident, GRC tracks regulatory notification clocks (GDPR 72 hours, HIPAA 60 days), keeps the decision log, coordinates with legal and insurers, and later updates the risk register and control set.

Why this matters: Showing how GRC adds value in a crisis (not just during audits) makes you stand out.

Q: Why are post-incident reviews (lessons learned) important, and how do you run one?

Intermediate

Purpose: understand what happened and why, and improve — not to assign blame.

- Hold it soon after the incident while details are fresh.
- Include all involved teams and relevant third parties.
- Build a timeline; identify what went well, what didn't, and the root causes (technical, process and people).
- Agree specific, owned, dated actions (e.g. new detections, playbook updates, training, control changes).
- Track actions to closure and feed findings into the risk register and future exercises.

Practical example: After a 9-hour outage, a blameless review finds an expired TLS certificate was the root cause. The actions are automated certificate monitoring, an owner for each certificate, and a new KRI: 'certificates expiring in under 30 days'.

Why this matters: Organisations that learn from incidents get breached less often and recover faster the next time.

Q: What is a crisis management team, and how is it different from the incident response team?

Intermediate

- The incident response team (technical) handles detection, containment, eradication and recovery of the incident.
- The crisis management team (executive) handles the strategic impact: business decisions, stakeholder and media communication, legal and regulatory matters, customer impact and resources.
- They connect through a defined escalation path and a regular briefing rhythm.
- Not every incident is a crisis. Activation criteria decide when the crisis team is convened.

Practical example: During a ransomware attack, the IR team works on containment while the crisis team (CEO, COO, legal, communications, CISO) meets every 4 hours to decide on clinic operations, patient communication and regulator notifications.

Why this matters: Mixing technical and executive roles slows both down. Clear separation with good links speeds up decisions.

Q: What should an incident response plan and playbooks include?

Intermediate

- **Plan:** purpose, scope, definitions and severity levels, roles (incident commander, technical lead, communications, legal, GRC), escalation, communication rules and evidence handling.
- **Playbooks:** step-by-step guides for common incident types, such as ransomware, BEC, data leakage, lost devices, DDoS and insider threats.
- **Contacts:** internal, external (forensics retainer, insurer, legal counsel, regulators, law enforcement).
- **Notification matrix:** which laws and contracts require notification, to whom and by when.
- Post-incident review and plan maintenance.

Practical example: The BEC playbook lists exactly what to do in the first hour: contact the bank's fraud desk, preserve the email headers, reset the mailbox password, remove forwarding rules and inform the finance director.

Why this matters: A good plan and practical playbooks turn panic into a routine. They are also among the first things insurers and auditors ask to see.

Q: What is digital forensics readiness, and why does it matter for GRC?

Intermediate

Forensic readiness means preparing in advance so that evidence can be collected and used effectively after an incident:

- Keep enough logs, for long enough, and protect their integrity (centralised, time-synchronised, tamper-resistant).
- Define how evidence is preserved, with chain of custody.
- Have a forensics retainer, ideally pre-approved by the insurer.
- Train first responders not to destroy evidence (for example, by wiping and rebuilding a machine too early).

Practical example: Because firewall and EDR logs were kept for 12 months, investigators could show attackers accessed only a test server, not the patient database, and notification was not required.

Why this matters: Without evidence you cannot confirm what data was accessed, which can force wider notifications and weaken your legal position.

Q: Your BIA says the payroll system has an RTO of 4 hours, but IT says recovery takes 3 days. What do you do?

Intermediate · Scenario

- **Validate both sides:** is 4 hours a real business need (payroll runs monthly, so maybe 48 hours is fine near the payroll date), and is IT's 3 days based on a recent test?
- **Discuss with the process owner what happens at each point:** 4 hours, 1 day, 3 days, especially near payroll dates.
- **If the gap is real, present options with costs:** faster recovery architecture, manual workarounds (re-run the previous month's payroll), or accepting the risk.
- Record the agreed RTO and the decision, and retest.

Practical example: The process owner agrees that a 48-hour RTO is enough except in the last 3 days of the month, when a manual fallback (repeat last month's payroll and adjust later) is documented and tested.

Why this matters: RTO gaps are common. Resolving them with facts and options, rather than just raising a finding, is what makes BCM credible.

Q: An employee's laptop with customer data is stolen from a car. Walk through the response.

Intermediate · Scenario

1. **Report quickly:** the employee reports to IT/security and police; get a police reference.
2. **Check the device status:** was it encrypted, was MFA required, when did it last check in? Trigger remote lock or wipe.
3. Revoke sessions and reset credentials stored on the device.
4. **Assess the data:** what was stored locally, how sensitive, how many people.
5. Decide on notification. If the device was fully encrypted and the key was not compromised, many laws treat the data as protected (for example HIPAA's encryption safe harbour), but document the assessment.

6. **Learn:** is local storage of sensitive data necessary? Are encryption and MDM compliance enforced on every device?

Practical example: MDM confirms the laptop was encrypted with BitLocker and last checked in that morning. It is remotely wiped when it next connects, sessions are revoked, and the incident is closed as 'no breach' with documented reasoning.

Why this matters: Lost-device incidents are frequent and simple if the basics (encryption, MDM, cloud storage) are in place. They become serious if not.

Q: A DR test fails halfway through. How do you handle it?

Intermediate · Scenario

- First, make sure production is safe and restore normal operations if the test affected them.
- Record exactly what failed, when and why. A failed test is valuable information, not a disaster.
- Run a short lessons-learned session with IT and business owners.
- Create corrective actions with owners and dates, and update the risk register (recovery capability is weaker than assumed).
- Schedule a retest, and inform governance of the result and the plan.
- Do not report it as a 'partial success' to leadership. Be honest about what it means for RTO.

Practical example: The database failover worked, but the application could not connect because DNS records were hard-coded. The fix is automated DNS failover, and a retest a month later meets the RTO.

Why this matters: The purpose of testing is to find failures before a real disaster does. Hiding failed tests defeats the purpose.

Q: A key supplier for clinical consumables goes out of business suddenly. How does BCM help?

Intermediate · Scenario

- The BIA and supplier mapping should already identify critical suppliers and single points of failure.
- **Activate the continuity plan for that dependency:** alternate suppliers, emergency stock, mutual aid with other facilities.
- Communicate with clinical leaders about what is affected and for how long.
- Afterwards, review supplier diversification, stock levels and financial monitoring of critical suppliers (a TPRM link).

BCM is not only about IT. Supply chain disruptions are a common real-world continuity event.

Practical example: The hospital switches to a pre-approved second supplier within 48 hours, using its 2-week safety stock in the meantime. Financial health checks on Tier 1 suppliers are added to the TPRM programme.

Why this matters: Interviewers check that you see continuity as a business discipline, not just IT disaster recovery.

Q: How do you align business continuity with ISO 22301 and ISO 27001 control 5.30?

Advanced

- ISO 22301 provides the full business continuity management system: context, BIA and risk assessment, strategies, plans, exercising, evaluation and improvement.
- ISO 27001:2022 control 5.29 covers information security during disruption, and 5.30 (ICT readiness for business continuity) requires ICT readiness to be planned, implemented, maintained and tested based on business continuity objectives.

- Practically, the BIA sets RTO and RPO, and IT must prove its DR capability meets them through testing.
- One integrated programme avoids two inconsistent sets of recovery targets.

Practical example: The BIA requires the EHR to be back in 4 hours, but the last DR test took 9 hours. Control 5.30 evidence shows the gap, and a project to automate failover is funded, with a retest that achieves 3.5 hours.

Why this matters: Auditors look for the link between business requirements (BIA) and proven ICT capability. A gap between them is a common finding.

Q: Should an organisation pay a ransom? What are the GRC considerations?

Advanced

This is a business and legal decision for leadership, not a GRC or IT decision, but GRC prepares the decision framework:

- **Legal:** sanctions rules (paying a sanctioned group may be illegal), local laws and law enforcement guidance.
- **Practical:** payment does not guarantee decryption or deletion of stolen data, and may mark you as a repeat target.
- **Insurance:** policy terms, insurer involvement and approval.
- **Recovery alternatives:** are backups intact and how long will restore take?
- Patient safety and critical service impact.

Decide the principles in advance, in calm conditions, and document the decision during the incident.

Practical example: During a tabletop, the board agrees a policy: no payment unless backups have failed and there is a risk to patient safety, and then only after legal sanctions checks and insurer involvement. In the real incident, backups worked and no payment was considered.

Why this matters: Ransom decisions made under pressure without a framework are often poor. Preparing the criteria in advance is a governance responsibility.

Q: You are the GRC lead in the first hour of a ransomware attack on a hospital. What are your priorities?

Advanced · Scenario

- Support the incident commander. Do not run the technical response yourself.
- Start the decision log and the regulatory clock tracker (which laws, which authorities, which deadlines).
- **Make sure patient safety is escalated:** clinical leaders activate downtime procedures.
- Bring in legal counsel, the insurer and the forensics retainer through the agreed channels.
- Prepare short, factual updates for executives and the board at a set rhythm.
- Ensure evidence preservation is part of containment decisions.
- Track third parties affected or involved (vendors, HIE connections, customers).

Practical example: Within the first hour, the GRC lead has a decision log running, has confirmed that the health regulator and the cyber insurer must be informed, has set 2-hourly executive updates, and has listed 3 vendors with live connections that must be checked.

Why this matters: GRC's value in a crisis is structure: decisions recorded, obligations met and leadership informed, so technical teams can focus on recovery.

Q: A cyber incident is taking longer than expected, and leadership wants to announce it is 'resolved'. How do you advise them?

Advanced · Scenario

- **Explain the risk of announcing too early:** if new facts appear (more data affected, attackers still present), the organisation loses credibility and may face regulatory criticism.
- **Recommend factual, accurate wording:** what is known, what is being done, and when the next update will be given.
- Check what notifications have been made and whether the public statement is consistent with them.
- **Define 'resolved' with the IR team:** containment confirmed, eradication validated, monitoring in place.
- Involve legal and communications in final wording.

Practical example: Instead of 'the incident is resolved', the statement says 'systems are restored and operating normally; our investigation with independent experts continues, and we will update affected patients directly if needed'. Two weeks later, new findings require a small number of extra notifications, and the earlier statement remains accurate.

Why this matters: Accurate communication during incidents protects trust. Premature 'all clear' messages are a common and damaging mistake.