TOPIC ONE

1

DOMAIN

Security Principles

26% of the exam

The largest domain on the exam, and the vocabulary the rest is built on

CANADA AI
& COMPLIANCE INC.

Why this domain gets three days

26%

of the exam — the largest single domain

9

lessons in this topic

~90

minutes

The gap this course fills

- The adaptive portal moves fast through what you already know
- 100% progress means the engine is satisfied — not that you would pass
- Ninety minutes is not enough for a quarter of the exam

So the method changes

- Portal is the first pass; these slides are the second
- Answer practice questions from day one, not at the end
- Keep an error log: question, why missed, lesson number

LESSON MAP

The nine lessons

1. **The security triad**
   Confidentiality, integrity, availability — and the trade-offs between them
2. **Identity and authentication**
   IAAA, the three factor types, multi-factor, non-repudiation
3. **Privacy**
   Privacy versus confidentiality, data roles, PII and PHI
4. **Risk foundations**
   Vocabulary, identification, qualitative and quantitative assessment
5. **Risk treatment**
   The four treatments, tolerance, residual risk, who may accept
6. **Security controls**
   Technical, administrative, physical — and how to sort them
7. **Governance**
   Laws, standards, policies, procedures, guidelines, baselines
8. **Professional ethics**
   The canons and their order of precedence
9. **Consolidation**
   Terminology, flashcards, and the memorisation set

Lessons 1 and 2 carry the vocabulary reused across every other domain. Do not rush them.

HOW TO USE THIS DECK

Two habits, starting now

Keep an error log

- One line per missed question
- Record the question, why you missed it, and the lesson number
- You will re-read this on your final day
- It will be worth more than any re-reading of the material

Say definitions out loud

- In your own words, not the book's
- Recognition is not enough, distractors are built to defeat it
- If you can produce a definition, you will not lose the mark
- If you can only recognise it, you probably will

The failure mode in this domain

- A comfortable feeling of recognition that collapses the moment you have to choose between four similar options

1.1

LESSON

The Security Triad

OBJECTIVE

Name the three properties security exists to protect, identify which one an incident violated, and explain why strengthening one can weaken another.

- Confidentiality, integrity and availability defined
- What breaks each, including accidents and hardware faults
- The three-question test for classifying an incident
- Trade-offs, and why security is not "maximise all three"

LESSON 1.1

Nearly all of security reduces to three properties

Confidentiality, integrity and availability.
No relation to the agency.

- These three are the properties that security controls exist to protect
- Every control you will meet in this course serves at least one of them
- Every incident you will be asked about violated at least one of them
- Get these three fluent and a surprising share of the exam becomes navigable

LESSON 1.1 — CONFIDENTIALITY

Confidentiality

Information is available only to those authorised to see it.
A restriction on WHO CAN READ.

- Fails when someone who should not have seen data sees it
- A mis-sent email · a stolen laptop · an exposed database
- An employee browsing records they have no business in
- Typical controls: encryption, access control, classification, need-to-know

LESSON 1.1 — INTEGRITY

Integrity

Information is accurate, complete, and altered only by authorised parties
in authorised ways. A restriction on WHO CAN CHANGE.

- Fails when data is modified without authorisation — a tampered transfer amount, a defaced page
- ALSO covers unauthorised deletion, not only alteration
- ALSO covers accidental corruption — a failing disk flipping bits is an integrity problem
- Typical controls: hashing, digital signatures, change control, input validation

LESSON 1.1 — AVAILABILITY

Availability

Information and systems are accessible to authorised users when needed.
A restriction on WHO CAN BE DENIED.

- Fails when legitimate users cannot reach what they are entitled to
- A denial-of-service attack · ransomware encrypting production data
- A severed fibre line · a server room losing power · a failed backup restore
- Note: no attacker is required. Hardware, human error and weather all break availability

The three properties together

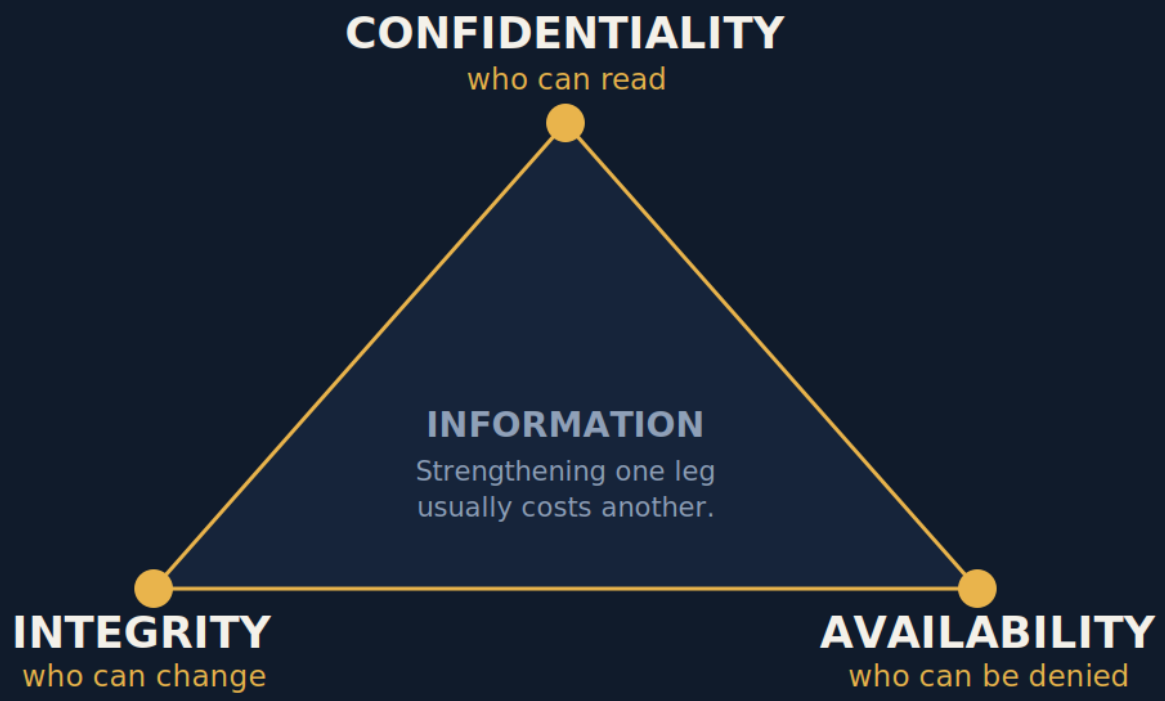

- Most real incidents break more than one property — ransomware breaks availability, and confidentiality too if data is published

When a question asks for one property, it wants the primary or first failure. Read for that word.

LESSON 1.1 — THE TEST

One question answers most triad items

Was something READ, CHANGED, or DENIED?

- READ → confidentiality
- CHANGED → integrity (including deleted, or accidentally corrupted)
- DENIED → availability
- Apply it before looking at the answer options. The options are designed to be plausible

LESSON 1.1 — TRADE-OFFS

Every control costs something elsewhere

Encrypt everything

- Protects confidentiality
- Lose the key and availability is gone — permanently
- A single administrator holding the only key is an availability risk, not a confidentiality one

Require multi-factor

- Protects confidentiality
- A user with a dead phone cannot work
- That is an availability cost, and a real one for field staff

Strict change control

- Protects integrity
- Slows emergency fixes
- An availability cost that has to be justified against the integrity gain

LESSON 1.1 — BALANCE

Security is not "maximise all three"

It is choosing a balance appropriate to the asset.

- A public marketing website needs availability far more than confidentiality
- A list of undercover informants needs confidentiality so badly that you accept it being hard to reach
- Medical records need all three — which is why healthcare security is expensive
- The classification of the asset drives which leg you favour

LESSON 1.1 — WORKED EXAMPLES

Which properties failed?

A ransomware attack encrypts a hospital's records; attackers also publish a sample online.

Availability (staff cannot reach records) and confidentiality (data published). A question asking for the primary impact of ransomware wants availability.

An employee changes the sales figures in a report before it reaches the board.

Integrity only. Nothing was disclosed to an outsider and nobody was denied access.

A backhoe cuts the fibre line to a company's data centre.

Availability only. No attacker, no disclosure — a reminder that failures are not always attacks.

A DNS poisoning attack redirects users to a convincing fake banking site.

Integrity failed first (DNS records altered), then confidentiality as users hand over credentials. One failure cascading into another.

LESSON 1.1 — QUICK CHECK

Answer before advancing

1. A web server is overwhelmed by botnet traffic and customers cannot place orders. Which property is primarily affected?

2. A hospital protects records from outsiders but sells identifiable patient data to advertisers without consent. Has confidentiality been violated? Has privacy?

3. Give one example of a control that improves integrity but costs availability.

Do not read past this slide. The whole point of the check is that you attempt it cold.

LESSON 1.1 — ANSWERS

Check your reasoning

1. Availability

No data was read or altered; authorised users were denied access. This describes a denial-of-service attack.

2. Confidentiality no, privacy yes

The holder controlled access and disclosed deliberately through an authorised channel, so confidentiality held in the technical sense. The data subjects lost control and never consented, so privacy failed. Lesson 1.3 develops this.

3. Several valid answers

Strict change control or approval workflows; input validation that rejects legitimate but unusual data; file integrity monitoring that quarantines on suspicion; mandatory code review before deployment.

1.2

LESSON

Identity and Authentication

OBJECTIVE

Separate identification, authentication, authorisation and accounting; classify authentication factors; explain what makes authentication multi-factor; and state what non-repudiation requires.

- The four steps and why their order matters
- The three factor categories, plus the two supplementary ones
- Why two prompts is not the same as two factors
- Non-repudiation as an assembled property

LESSON 1.2 — THE FOUR STEPS

I claim, I prove, I'm permitted, I'm recorded

Being authenticated is not the same as being permitted.

Identity proofing comes before all four

- Before a credential is issued, the organisation should establish that you are a real person entitled to one — checking a passport at onboarding, verifying employment

LESSON 1.2 — THE FOUR STEPS

Each step, precisely

Identification	You claim who you are. Typing a username, presenting a badge, stating an account number. An unverified assertion — anyone can claim to be anyone.
Authentication	You prove the claim. Supplying the password that goes with the username. The system verifies your evidence against what it holds on record.
Authorisation	The system determines what you, now authenticated, are permitted to do. Being authenticated does not mean being permitted.
Accounting	The system records what you did. Logs, audit trails, session records. What makes actions traceable to an identity afterwards.

Mnemonic: I claim, I prove, I'm permitted, I'm recorded.

LESSON 1.2 — FACTORS

Three types of evidence

Count categories, never prompts.

Something you KNOW

Password, PIN,
passphrase,
security question

Something you HAVE

Token, smart card,
phone receiving
a one-time code

Something you ARE

Fingerprint, iris,
face, voice
(biometrics)

Multi-factor = evidence from two DIFFERENT boxes

Password + security question = both KNOW = still single-factor.

Also exists: somewhere you are (location) and something you do (behaviour — typing rhythm, gait). Know they exist; the three above carry the exam.

LESSON 1.2 — THE MFA RULE

The rule people get wrong

Multi-factor requires evidence from two or more DIFFERENT categories.
Count categories, never prompts.

- Password + security question = NOT multi-factor. Both are knowledge
- Two passwords = NOT multi-factor
- Fingerprint + facial scan = NOT multi-factor. Both are inherence
- Password + phone code = multi-factor. Knowledge plus possession
- Smart card + PIN = multi-factor. Possession plus knowledge

LESSON 1.2 — TRADE-OFFS

What each factor type costs you

Knowledge

- Cheap, no hardware needed

- Reusable across sites
- Guessable and phishable
- Forgettable — help-desk cost

Possession

- Resists remote guessing

- Tokens get lost and stolen
- Expensive to issue and replace
- Logistics for remote staff

Inherence

- Cannot be forgotten
- Hard to share

- CANNOT BE REISSUED after compromise
- Privacy concerns; measurable error rates

Biometric error rates — false acceptance, false rejection, crossover error rate — are covered properly in Domain 3.

LESSON 1.2 — NON-REPUDIATION

Non-repudiation is assembled, not installed

What it means

- A party cannot credibly deny having taken an action
- Classic framing: "I never approved that transfer"
- Not a product you buy — a property you build
- Closely related to authenticity: assurance that data came from its claimed source

What it requires

- Strong authentication — ties the action to one identity
- Integrity — the record cannot be altered afterwards
- Accounting — a record exists at all
- Commonly digital signatures, binding a private key to a specific message

Why shared accounts destroy it

- If five administrators use one login, every action is attributable to five people — which is to say, to nobody. Hence "no shared accounts", and it returns in Domain 3

LESSON 1.2 — WORKED EXAMPLES

Classify each

A user enters a password, then approves a push notification on their registered phone.

Multi-factor — knowledge plus possession. Two categories, so it qualifies.

A banking app requires a PIN and your mother's maiden name.

Single-factor, despite two prompts. Both are knowledge. This is the trap.

A door requires a badge swipe and a fingerprint.

Multi-factor — possession plus inherence.

A log shows account jsmith deleted a record at 14:22. Four people use jsmith.

Accounting occurred, but non-repudiation fails — no individual can be held to the action.

LESSON 1.2 — QUICK CHECK

Answer before advancing

1. A system requires a password and a one-time code sent by SMS. Which categories are in use, and is this multi-factor?

2. A user logs in successfully but receives "access denied" opening the payroll folder. Which of the four steps refused them?

3. Which factor category cannot be replaced after compromise, and why does that matter?

LESSON 1.2 — ANSWERS

Check your reasoning

1. Knowledge and possession — yes, multi-factor

The password is knowledge, the phone receiving the code is possession. Worth noting SMS is the weakest common delivery method for a possession factor, because SIM-swapping lets an attacker receive the code without holding the phone. Still multi-factor, just not strong multi-factor.

2. Authorisation

Authentication succeeded — that is why they got in at all. Authorisation then determined they were not permitted this particular resource.

3. Inherence — biometrics

You can issue a new password or token. You cannot issue a new fingerprint or iris. A compromised biometric template is compromised permanently, which is why biometric data warrants unusually strong protection.

1.3

LESSON

Privacy

OBJECTIVE

Distinguish privacy from confidentiality, identify the roles in a privacy regime, and recognise the data categories that trigger extra handling.

- Why privacy and confidentiality are not synonyms
- The scenario the exam builds around that distinction
- Data subject, controller, processor, owner, custodian
- PII and PHI, and how much regulation detail you actually need

LESSON 1.3 — THE DISTINCTION

Privacy is not confidentiality

Confidentiality

- An obligation the HOLDER of data has
- Protect it from those not authorised to see it
- Fails when the wrong person reads something
- A duty owed by whoever holds the data

Privacy

- A right the SUBJECT of data has
- Control over information about themselves
- Fails when data is used in ways never agreed to
- A right held by the person the data describes

The consequence

- An organisation can keep your data perfectly confidential and still violate your privacy, by using it in ways you never consented to

LESSON 1.3 — THE SCENARIO

The exam's favourite privacy scenario

A hospital protects records perfectly from outsiders, then sells identifiable patient data to advertisers without consent.

- Confidentiality was NOT violated — access was controlled, and disclosure was deliberate and authorised by the holder
- Privacy WAS violated — the subjects lost control over information about themselves and never consented
- The technical control worked exactly as designed. The failure was one of permission, not protection
- If a question describes authorised-but-unwanted use of data, the answer is privacy

Privacy at work is a genuine tension

Employees have some expectation of privacy.
Employers have legitimate monitoring interests.

- Monitoring of email, browsing and systems use is common and often lawful
- What matters is whether the expectation was set — usually via an acceptable use policy signed at onboarding
- Personal device use blurs whose data is whose, which is why BYOD policy exists (Domain 5)
- The exam's position: monitoring is acceptable when disclosed, and problematic when covert

LESSON 1.3 — ROLES

Who is who in a privacy regime

Data subject

- The person the data is about
- Holds the privacy right
- Source of consent

Controller

- Decides why and how data is processed
- Carries the primary obligation

Processor

- Processes data on the controller's instructions
- Acts on behalf of, not independently

Owner / custodian

- Owner is accountable for the data
- Custodian maintains it day to day

Owner and custodian are internal roles; controller and processor are regulatory ones. They overlap but are not the same vocabulary.

LESSON 1.3 — CATEGORIES

Data that triggers extra handling

The two categories to know

- PII — personally identifiable information: anything that identifies a person, alone or combined with other data
- PHI — protected health information: medical and health-payment data, typically under its own stricter regime

Regulation: shape, not detail

- Know that regimes create rights for subjects and obligations for controllers
- Know that breach notification duties commonly exist
- The exam wants the concept, NOT article numbers or specific statutes
- Do not rabbit-hole into GDPR or HIPAA specifics

LESSON 1.3 — QUICK CHECK

Answer before advancing

1. A company shares customer data with a marketing partner, as permitted by terms customers accepted. Privacy issue, confidentiality issue, both, or neither?

2. An employer reads an employee's work email without notice or policy. Which concept is most at stake?

3. Who holds the primary obligation in a controller / processor relationship?

LESSON 1.3 — ANSWERS

Check your reasoning

1. Arguably neither, and that is the point

Confidentiality holds — disclosure was authorised. Privacy is respected in form because consent was obtained. Whether the consent was meaningful is an ethics and regulation question rather than a technical one, which is exactly where these debates actually sit.

2. Privacy

The employee lost control over information about themselves without having been told the monitoring would occur. Disclosed monitoring under an acceptable use policy would change the answer.

3. The controller

The controller decides the purpose and means of processing, so it carries the primary obligation. The processor acts on instruction, though it has duties of its own.

1.4

LESSON

Risk Management Foundations

OBJECTIVE

Use risk vocabulary precisely, describe where risk identification inputs come from, and choose between qualitative and quantitative assessment.

- Why organisations manage risk formally at all
- Asset, threat, threat actor, vulnerability, likelihood, impact, risk
- Why a vulnerability alone is not a risk
- Qualitative versus quantitative assessment, and when each fits

LESSON 1.4 — WHY BOTHER

Why formalise risk at all?

Because "react to whatever is loudest" does not scale, and cannot be defended to a board.

- Prioritisation — you cannot protect everything equally, so something must decide the order
- Budget justification — a documented risk is an argument for spending; a hunch is not
- Accountability — a named risk has a named owner and a recorded decision
- Consistency — two similar risks get similar treatment across the organisation

LESSON 1.4 — VOCABULARY

The core terms, strictly

Asset	Something of value worth protecting — data, systems, people, facilities, reputation.
Threat	A potential cause of an unwanted incident. A possibility, not an event.
Threat actor	The person or group behind a threat. Insider, criminal, competitor, state, or accidental actor.
Vulnerability	A weakness that a threat could exploit. Unpatched software, an untrained user, an unlocked door.

LESSON 1.4 — VOCABULARY

Then the three that combine

Likelihood	The chance of the unwanted thing actually happening. Often estimated rather than calculated.
Impact	What it costs if it does happen — financial, operational, reputational, regulatory, human.
Risk	The combination of likelihood and impact. Conceptually likelihood × impact, though rarely a literal calculation.
Threat vector	The path or means by which a threat reaches the vulnerability — email, removable media, a supplier.

LESSON 1.4 — THE KEY INSIGHT

Both halves are required

A vulnerability with no threat is not a risk.
A threat with no vulnerability is not a risk.

- An unpatched server on a network with no attackers and no connectivity presents little risk
- A determined attacker facing a fully patched, segmented system presents little risk
- Risk exists where a threat can reach a vulnerability that affects an asset of value
- This is why controls work: they remove the vulnerability, block the vector, or reduce the impact

LESSON 1.4 — IDENTIFICATION

Where the inputs come from

Sources for risk identification

- Asset inventory — you cannot assess what you have not listed
- Vulnerability scans and penetration test results
- Audit findings and prior incident history
- Staff and process owners, who know where the workarounds are
- Threat intelligence and industry reporting

Why the inventory comes first

- Unknown assets cannot be protected, assessed, patched or monitored
- "Shadow IT" is the standard example — systems the security team does not know exist
- Asset inventory is the unglamorous foundation of every later step

LESSON 1.4 — ASSESSMENT

Two ways to assess, both valid

Qualitative

- Ordinal scales — high / medium / low, or 1 to 5
- Heat maps plotting likelihood against impact
- Fast, cheap, and workable with no historical data
- Subjective, and hard to compare across teams
- The default in most organisations

Quantitative

- Expressed in currency — expected loss per year
- Requires credible data on frequency and cost
- Comparable, and directly usable in budget arguments
- Expensive, and false precision is a real hazard
- Used where good data exists — fraud, hardware failure

Choosing between them

- Fit the method to the data you actually have. Quantitative assessment built on invented numbers is worse than an honest qualitative one

LESSON 1.4 — QUICK CHECK

Answer before advancing

1. A server has an unpatched vulnerability but is powered off and disconnected. Is there a risk? Explain.

2. Which comes first: the asset inventory, or the vulnerability assessment? Why?

3. An organisation has no historical incident data. Which assessment method fits?

LESSON 1.4 — ANSWERS

Check your reasoning

1. Very little

The vulnerability exists, but with no connectivity there is no vector for a threat to reach it. Risk requires a threat that can actually reach a vulnerability affecting something of value. Note "very little" rather than "none" — physical access is still a vector.

2. The asset inventory

You cannot scan, assess or protect what you do not know exists. Every subsequent step depends on knowing the population of assets, which is why shadow IT is a security problem rather than merely an administrative one.

3. Qualitative

Quantitative assessment needs credible frequency and cost data. Without it, the numbers would be invented, and an invented figure carries false authority that a high/medium/low rating does not.

1.5

LESSON

Risk Treatment and Tolerance

OBJECTIVE

Match a described response to one of the four treatments, distinguish inherent from residual risk, and state who may accept a risk.

- The four treatments as a closed set
- Why transference never moves the whole risk
- Inherent, residual and the role of tolerance
- Who has authority to accept risk — and who does not

LESSON 1.5 — THE PROCESS

The process, and the four treatments

FOUR TREATMENTS — the whole set

AVOID	MITIGATE	TRANSFER	ACCEPT
stop the activity	add controls	insure or contract out	document and live with it

What is left after treatment is residual risk. Only senior management accepts risk.

Assessment is qualitative or quantitative. Monitoring exists because risk is not static — new threats and new assets change the picture.

LESSON 1.5 — THE FOUR

The four treatments

Avoidance	Stop doing the activity that creates the risk. Discontinuing a product line, not entering a market, removing a feature. The only treatment that eliminates the risk entirely.
Mitigation	Add controls to reduce likelihood or impact. The most common treatment, and where most of security engineering lives.
Transference	Shift some consequence to another party — insurance, or a contractual clause pushing liability to a supplier.
Acceptance	Acknowledge the risk and proceed without further treatment. Legitimate when the risk sits within tolerance, and must be documented.

LESSON 1.5 — TRANSFERENCE

Transference never moves the whole risk

Insurance pays. It does not prevent.

- A cyber insurance policy may cover the financial loss from a breach
- It does not restore the data, undo the disclosure, or repair the reputational damage
- Regulatory penalties are frequently uninsurable
- The operational disruption still happens, and you still have to respond to it
- So transference is usually paired with mitigation, not used instead of it

LESSON 1.5 — BEFORE AND AFTER

Inherent and residual risk

Inherent risk

- The risk level before any treatment is applied
- The starting point of the assessment
- Sometimes called gross risk

Residual risk

- What remains after treatment
- Never zero — controls reduce, they do not eliminate
- Must be compared against tolerance
- If residual risk still exceeds tolerance, more treatment is needed

The loop this creates

- Assess inherent risk → treat → measure residual risk → compare to tolerance → treat again or accept. Acceptance is the exit from the loop, not a step within it

LESSON 1.5 — TOLERANCE

Tolerance is what makes the decision

Risk tolerance is set at senior level and expressed as a threshold. It decides whether a given risk gets treated or accepted.

- Two identical risks at two organisations can get opposite decisions, correctly
- A startup may accept risks a bank could not, because their tolerances differ
- Risk appetite is the broader statement of how much risk the organisation seeks; tolerance is the operational limit
- Neither is set by the security team — they are business decisions the security team informs

LESSON 1.5 — AUTHORITY

Who is allowed to accept a risk

Senior management, or the designated risk owner.
Never the practitioner who found it.

- The person who identifies a risk documents it and escalates it — they do not decide to live with it
- This is a separation-of-duties principle: finding and accepting are different roles
- It appears in ethics scenarios too: being told to "just accept" a finding you have no authority to accept
- And in incident response: accepting residual risk after recovery is a management decision

LESSON 1.5 — WORKED EXAMPLES

Name the treatment

A company discontinues a payment feature after finding it cannot be secured economically.

Avoidance. The activity creating the risk was stopped entirely.

A company buys cyber insurance covering breach response costs.

Transference — partial. The financial consequence shifts; the operational and reputational consequences do not.

A company deploys MFA across all remote access.

Mitigation. A control reducing the likelihood of credential compromise.

A company documents that a legacy system cannot be patched, and the CIO signs off on continuing to run it.

Acceptance — done correctly, because it is documented and signed by someone with authority.

LESSON 1.5 — QUICK CHECK

Answer before advancing

1. A team lead discovers a serious vulnerability and decides the business can live with it. What is wrong here?

2. Why is buying insurance rarely a complete response to a security risk?

3. Two companies face the same risk. One treats it, one accepts it. Can both be correct?

LESSON 1.5 — ANSWERS

Check your reasoning

1. The team lead almost certainly lacks the authority

Identifying a risk and accepting it are separate roles. The finder documents and escalates; senior management or the named risk owner decides. Accepting a risk you found is a control failure regardless of whether the judgement was sound.

2. Because it transfers financial consequence only

Insurance does not restore data, undo disclosure, repair reputation, or prevent operational disruption, and regulatory penalties are often uninsurable. It is normally paired with mitigation rather than used instead of it.

3. Yes

Risk tolerance is set by each organisation's senior management and legitimately differs. The same residual risk can sit inside one organisation's tolerance and outside another's. The decision is only wrong if it was made by someone without authority, or without documentation.

1.6

LESSON

Security Controls

OBJECTIVE

Sort any described control into technical, administrative or physical, and explain how controls connect back to risk.

- The three categories and what enforces each
- The sorting rule for ambiguous cases
- A practice set to run cold
- Why a control exists at all — the link back to tolerance

LESSON 1.6 — THE CATEGORIES

Three categories, one sorting rule

Classify by the mechanism doing the work — not by the goal.

TECHNICAL

Enforced by technology

Firewalls, encryption, access control lists, antivirus, MFA

ADMINISTRATIVE

Enforced by process

Policies, training, background checks, access reviews

PHYSICAL

Enforced by barriers

Locks, fences, guards, badge readers, cameras

A control is a mitigation, chosen because a risk exceeded tolerance.

Some controls arguably span categories. Pick the best answer by asking what mechanism actually does the work.

LESSON 1.6 — THE RULE

Classify by mechanism, not by goal

Ask: what is actually doing the work?

- A badge reader is PHYSICAL — the door and the lock are doing the work, even though electronics are involved
- An access control list is TECHNICAL — software is enforcing it
- A policy requiring badge use is ADMINISTRATIVE — a process, enforced by people
- All three might serve the same goal of controlling entry. The goal does not determine the category

LESSON 1.6 — PRACTICE

Sort these before you look

Firewall rule set

Technical

Encryption at rest

Technical

Employee background check

Administrative

Perimeter fence

Physical

Locked server cabinet

Physical

Quarterly access review

Administrative

Security awareness training

Administrative

Intrusion detection system

Technical

Cover the amber answers, work down the list, then check. Anything missed goes in the error log with its lesson number.

LESSON 1.6 — PRACTICE

A harder set

Badge reader on a door

Physical — the door does the work

Incident response plan

Administrative

Mandatory vacation policy

Administrative

Bollards outside a building

Physical

CCTV camera

Physical

Password complexity enforcement in software

Technical

Automatic screen lock after 5 minutes

Technical

Written password policy

Administrative

Note the last two: the same subject, different categories, because different mechanisms enforce them.

LESSON 1.6 — THE LINK BACK

Why any control exists

A control is a mitigation, chosen because a risk exceeded tolerance.

- Controls are not selected because they are available or fashionable
- The chain: asset → threat → vulnerability → risk → exceeds tolerance → treat by mitigating → select control
- This is why "we bought a tool" is not a security strategy
- And why questions about control selection usually have a risk-based answer

1.7

LESSON

Governance

OBJECTIVE

Rank the governance elements by where their authority originates, and identify which element a one-sentence description refers to.

- Laws and regulations, standards, policies, procedures, guidelines
- Baselines, and what makes them different
- Why the hierarchy matters operationally
- Discriminating the elements from a single sentence

The governance hierarchy

Authority flows downward. Nothing below may contradict anything above.

LAWS & REGULATIONS
External. Mandatory. Legal penalty for breach.

STANDARDS
Specify how something must be done.

POLICIES
Internal management direction. Mandatory inside the org.

PROCEDURES
Step-by-step instructions to carry out policy.

GUIDELINES
Advisory. Recommended, not required.

Baselines set the minimum acceptable configuration. A procedure contradicting policy is a finding; a policy contradicting regulation is invalid.

LESSON 1.7 — THE ELEMENTS

Each element, and its source of authority

Laws & regulations	Imposed externally by government or regulator. Mandatory, with legal penalty for breach. The organisation does not get to choose.
Standards	Specify how something must be done — a required configuration, protocol or method. May be external (industry) or internal.
Policies	Internal management direction. Mandatory inside the organisation. State what must happen and why, not how.
Procedures	Step-by-step instructions for carrying out a policy. The operational detail — who does what, in what order.
Guidelines	Advisory. Recommended practice, not required. The element that is explicitly optional.

Baselines, and why the hierarchy matters

Baselines

- The minimum acceptable level of security for a system or class of system
- A starting configuration everything must at least meet
- Deviations require justification and approval
- Returns in Domain 5 under system hardening and configuration management

Why the hierarchy is operational, not decorative

- A policy that contradicts a regulation is invalid — the regulation wins
- A procedure that contradicts policy is an audit finding
- A guideline that contradicts a policy is simply ignored
- Authority flows downward, and nothing below may override anything above

LESSON 1.7 — WORKED EXAMPLES

Name the element

"All customer data must be encrypted at rest."

A policy — internal management direction, mandatory, states what rather than how.

"Use AES-256 in GCM mode with keys rotated every 90 days."

A standard — it specifies how the policy is to be implemented.

"To encrypt a new database: open the console, select the instance, choose Encryption, then..."

A procedure — step-by-step operational instructions.

"Organisations processing health data must notify affected individuals within 60 days of a breach."

A regulation — externally imposed, mandatory, carries legal penalty.

LESSON 1.7 — QUICK CHECK

Answer before advancing

1. Which governance element is explicitly advisory rather than mandatory?

2. A team's documented procedure conflicts with company policy. What is this, and which wins?

3. Where do baselines sit, and what happens when a system cannot meet one?

LESSON 1.7 — ANSWERS

Check your reasoning

1. Guidelines

They are recommended practice, not requirements. Everything else in the hierarchy is mandatory within its scope — laws and regulations externally, standards, policies and procedures internally.

2. An audit finding, and the policy wins

Procedures implement policy, so a procedure cannot override it. The correct resolution is to fix the procedure — or, if the procedure reflects operational reality, to formally change the policy. What is not acceptable is leaving the conflict in place.

3. Below standards, as the minimum acceptable configuration

When a system cannot meet a baseline, the deviation is documented, justified, approved by someone with authority, and usually paired with a compensating control. An undocumented deviation is a finding.

1.8

LESSON

Professional Ethics

OBJECTIVE

Reason through an ethical conflict using canon precedence, and recognise the pattern the exam's ethics questions follow.

- Why canon order is precedence order
- The general pattern in ethics scenarios
- Worked conflicts
- Why you must read the official wording yourself

Applying the canons

Canon order is precedence order

1. Society, the common good, public trust, infrastructure
2. Acting honourably, honestly, justly, responsibly, legally
3. Diligent and competent service to principals
4. Advancing and protecting the profession

Where two canons conflict, the earlier one wins. Paraphrased above — read the official wording on the ISC2 website.

The general pattern

- Society and public trust before duty to the employer
- Disclosure through proper channels before unilateral action
- Never conceal a finding you are asked to hide

LESSON 1.8 — THE RULE

Canon order is precedence order

Where two canons conflict, the earlier one wins.

- This is the mechanism the exam tests, more than the wording of any individual canon
- Protecting society and public trust outranks duty to your employer
- Duty to your employer outranks advancing the profession
- So a scenario pitting employer interest against public safety has a determined answer

LESSON 1.8 — SCENARIOS

Worked conflicts

You discover your employer has suffered a breach affecting customers and has not disclosed it.

Public trust and the common good outrank duty to the employer. Raise it internally through proper channels first; escalate if suppressed.

A manager asks you to omit a serious finding from an audit report.

Refuse. Acting honestly and legally, and protecting public trust, both outrank the instruction. Concealment is not available.

A colleague is misrepresenting credentials they do not hold.

Protecting the profession applies, and so does acting honourably. Report through the appropriate channel rather than confronting publicly.

A client pressures you to skip a required control to hit a deadline.

Diligent and competent service does not mean compliance with unsafe instructions. Document the risk, escalate, and let an authorised person accept it.

LESSON 1.8 — THE SOURCE

Read the source, not a summary

Why this deck paraphrases

- This is an independent course and does not reproduce ISC2's text
- The paraphrase teaches the mechanism, which is what most questions test

Why you must still read the original

- The exam may echo ISC2's own phrasing
- The ordering must be exact, and ordering is the tested mechanism
- It is short — read the preamble and four canons directly on the ISC2 website before exam day

Do this on Day 3

- It takes five minutes and it is the one piece of Domain 1 you should get from the source rather than from any course

1.9

LESSON

Consolidation

OBJECTIVE

Produce the memorisation set from memory, recognise the exam's question patterns, and identify your own weak areas before moving on.

- The pure-memorisation set
- Six question patterns that cover most of the domain
- The traps that cost marks
- A three-day plan

LESSON 1.9 — MEMORISE

Produce these, don't just recognise them

Closed sets

- Four risk treatments:
- avoid · mitigate · transfer · accept

- Three control categories:
- technical · administrative · physical

- Three authentication factors:
- know · have · are

Ordered sets

- The triad:
- confidentiality · integrity · availability

- IAAA:
- identify · authenticate · authorise · account

- Governance:
- laws · standards · policies · procedures · guidelines

Rules and asymmetries

- Earlier canon wins

- Only senior management accepts risk

- Biometrics cannot be reissued

- Count categories, not prompts

- Classify by mechanism, not goal

Build these as flashcards before starting Domain 2. This vocabulary is reused in Domains 2 and 5.

EXAM TECHNIQUE

Six patterns cover most of this domain

1	**Scenario → property**	"An attacker modifies records. Which principle is violated?" Ask: read, changed, or denied.
2	**Property → control**	"Which control best protects confidentiality?" Encryption and access control — not backups.
3	**Classify the example**	Sort a control into technical, administrative or physical by its mechanism.
4	**Name the element**	Given one sentence, decide whether it is a policy, procedure, standard, guideline or law.
5	**Count the factors**	Given a login process, is it single- or multi-factor? Count categories.
6	**Apply the canons**	Given an ethical conflict, choose the earlier canon.

In practice sets, name the pattern before answering. It is faster and it exposes which pattern you are weak on.

COMMON TRAPS

Where marks get lost

Confidentiality ≠ privacy

- Holder's obligation versus subject's right
- Tested directly, and the most commonly missed distinction in the domain

Two prompts ≠ two factors

- Password plus security question is still single-factor
- Count categories, never prompts

Authentication ≠ authorisation

- Logged in but refused access is an authorisation question
- Authentication already succeeded

Availability is not minor

- A full third of the triad
- Where DoS, ransomware and redundancy questions live

And: assuming one incident breaks one property. Most break several — when a question wants one, it wants the primary or first failure.

Next: Domain 2

Incident Response, Continuity and Recovery

10% of the exam. Three distinct plans in a sequence, and the availability leg of the triad applied to systems and people.

Before you move on: can you produce — not just recognise — the four risk treatments, the three control categories, and the difference between privacy and confidentiality?