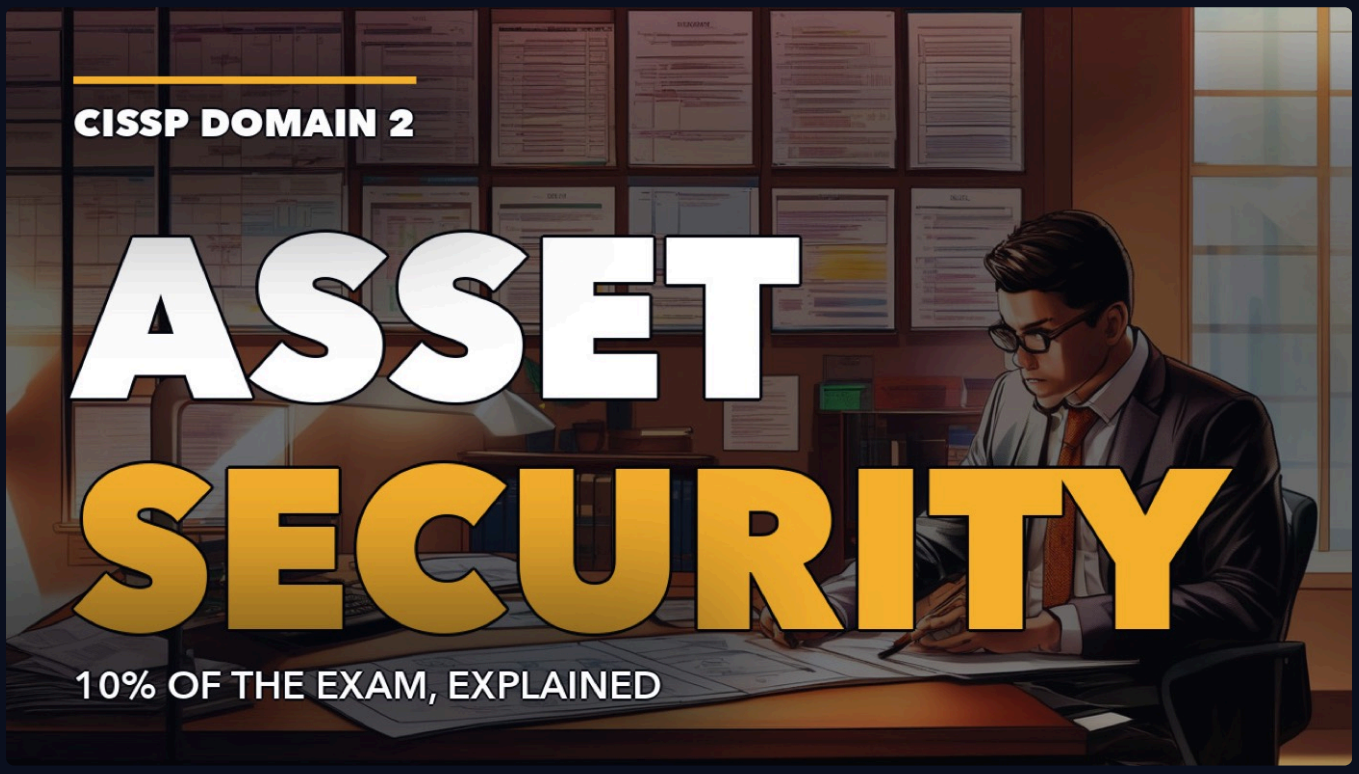
Forwards or Backwards

https://forwardorbackwards.com



CISSP Domain 2: Asset Security

Watch on YouTube: https://www.youtube.com/watch?v=W4r6Qm2S9vs

A full tour of Domain 2 of the CISSP exam - Asset Security - worth 10% of the exam under ISC2's exam outline effective 15 April 2024.

This video works through all six of Domain 2's sub-domains in order: identifying and classifying information and assets, with government and commercial classification schemes and a worked HR/finance example; extending classification to hardware, software and facilities through an asset register, with labelling and marking; handling requirements for storage, transport and destruction; secure provisioning and asset inventory and configuration management; the data roles - owner, custodian, user, subject, and GDPR's controller and processor; the full data lifecycle from collection and location through maintenance, retention, remanence and sanitisation, including NIST SP 800-88's Clear, Purge and Destroy categories and the 2025 update (Revision 2) that points to IEEE 2883 for technique selection; asset retention and the end-of-life/end-of-support distinction; the three data

states - at rest, in transit, in use; scoping and tailoring security baselines; selecting data security standards such as PCI DSS v4.0.1; data protection methods including DRM, DLP, CASB, tokenisation, masking, pseudonymisation and anonymisation; and privacy obligations, chiefly GDPR, running through the whole lifecycle.

It closes with how today's Domain 2 absorbed material from two domains in the pre-2015 ten-domain CBK, a full recap, and seventeen exam-style questions, each with four options, a pause to think, and an explained answer, ending with how Domain 2 connects to the other seven CISSP domains.

Every figure is on screen with its source and date.

Educational documentary. Not financial or investment advice.

Chapters

1. Welcome, and what this domain weighs
2. 2.1 Classifying information
3. 2.1 Classifying assets and marking
4. 2.2 Handling requirements
5. 2.3 Provisioning resources securely
6. 2.3 Asset inventory and management
7. 2.4 Data roles across the lifecycle
8. 2.4 Data lifecycle: collection
9. 2.4 Data lifecycle: location
10. 2.4 Data lifecycle: maintenance
11. 2.4 Data lifecycle: retention
12. 2.4 Data remanence
13. 2.4 Data sanitisation and destruction
14. 2.5 Asset retention, EOL and EOS
15. 2.6 Data states: at rest, in transit, in use
16. 2.6 Scoping and tailoring baselines
17. 2.6 Selecting data security standards
18. 2.6 Data protection: DRM, DLP, CASB
19. 2.6 Protection methods: tokenisation, masking, anonymisation
20. 2.6 Privacy obligations across the lifecycle
21. From the old 10-domain CBK to today's Domain 2
22. Domain 2 full recap
23. Review questions 1: classification, roles and handling
24. Review questions 2: lifecycle, retention and destruction
25. Review questions 3: controls, privacy, other domains

1. Welcome, and what this domain weighs





Welcome to this training module on Domain 2 of the CISSP exam: Asset Security.

Under ISC2's exam outline, Domain 2 is worth ten per cent of the exam. That outline has been in effect since 15 April 2024. The exam itself runs on a Computerized Adaptive Testing format, a hundred to a hundred and fifty questions, in three hours. The pass mark is seven hundred out of a thousand.

Ten per cent sounds modest next to the biggest domains, but asset security sits underneath almost everything else a security professional does. You cannot protect what you have not identified, and you cannot decide how hard to protect something until you know what it is worth.

This module takes every one of Domain 2's six sub-domains in order. Identifying and classifying information and assets. Establishing handling requirements. Provisioning resources securely. Managing the data lifecycle from collection to destruction. Ensuring appropriate retention. And determining data security controls and compliance requirements.

Picture the setting this whole module works in. A data centre humming with racked servers, a records room lined with boxes, and a security operations centre watching screens around the clock — much like the United States Department of Homeland Security runs as its National Operations Center. Every example in this module happens somewhere like that.

Each chapter defines its terms, works through at least one real-world scenario, then closes with key points and an exam tip: how ISC2 tends to test the idea, and the trap candidates fall into. At the end, a full recap, then exam-style review questions with explained answers.

Let's begin.

2. 2.1 Classifying information



Classification is the act of sorting information by how much harm its loss would cause, so that protection effort goes where it is actually needed.

Governments typically use a scheme such as Top Secret, Secret, Confidential and Unclassified, each level defined by the damage disclosure would cause to national security. Commercial organisations use their own equivalent scale, often Confidential, Private, Sensitive and Public, defined the same way but against business harm rather than national harm.

The criteria behind a classification decision are broader than secrecy alone. How much is the information worth? How old is it, and does it lose sensitivity with age? How long does it need to stay useful? And is it tied to a named person, which brings privacy law into play alongside business risk?

Picture a worked example. A mid-size company reviews its HR files and its finance files for the first time. Payroll records carrying bank details and salary figures are rated Confidential, the top commercial tier, because disclosure would damage employees directly. General finance reports used for the next board meeting are rated Private, protected but shared more widely inside the company. A public-facing price list is rated Public.

Two traps show up in exam questions on this topic. The first is treating classification as permanent: information can be reclassified as it ages or as circumstances change, and a scheme that cannot be revisited is not doing its job. The second is forgetting that over-classifying has a real cost too — it slows work down and trains people to ignore labels that are always set too high.

Key points: Classification sorts information by potential harm, not by instinct. Government and commercial schemes use different labels for the same idea. The criteria are value, age, useful life and personal association, not secrecy alone.

Exam tip: ISC2 tends to test whether you can apply a scheme consistently to a scenario, not whether you can recite the four government labels. Expect a short case and a question asking which level fits – think like the manager who has to defend the choice later.

3. 2.1 Classifying assets and marking



Classification does not stop at data. The same logic extends to the assets that store, process and carry information. Sort them by the harm their loss or compromise would cause — servers, laptops, backup tapes, even the buildings and facilities that house them.

An asset register lists every asset an organisation owns or is responsible for, together with its classification, its owner and its location. Hardware, software and facilities each need their own entry, because a facility housing Confidential-rated servers inherits protection requirements the building itself never had on its own.

Once an asset or a piece of information has a classification, that classification has to be visible to the people who handle it. **Labelling** marks the classification on the information or the asset itself — a header on a document, a sticker on a drive. **Marking** goes further, often adding handling instructions alongside the label: who may read it, how it must be stored, how it must be destroyed.

Consider a worked example. An audit of a finance department's laptop fleet finds three machines holding Confidential payroll exports with no classification sticker and no entry in the asset register at all. The fix is not just adding the sticker — it is adding the laptops to the register, assigning an owner, and setting a review date.

A common exam trap is assuming a digital file's classification travels with it automatically when it is copied or printed. It does not, unless the organisation's tooling and procedures make it travel. A label on screen is not the same as a label the next system, or the next person, can see and act on.

Key points: Classification extends from data to hardware, software and facilities. An asset register ties classification to an owner and a location. Labelling states the classification, marking adds handling instructions.

Exam tip: Expect a scenario testing whether you noticed an unregistered or unlabelled asset, not one testing the label names themselves. ISC2 wants you thinking about completeness of the register, which is where real breaches start.

4. 2.2 Handling requirements



Once something is classified, the organisation has to say, in writing, what that classification actually requires in practice. How it is stored. How it travels. Who may touch it. And how it is destroyed when its life ends.

A handling standard ties specific rules to each classification level. Confidential material might require encryption at rest, a locked cabinet or encrypted drive for storage, courier or encrypted transfer for transport, and two-person authorisation before destruction. Public material might require almost none of that.

Storage, transport and destruction are the three moments where handling requirements bite hardest. Each is a moment an asset leaves the direct control of the person responsible for it: sitting in a cabinet overnight, travelling between offices, or being thrown away.

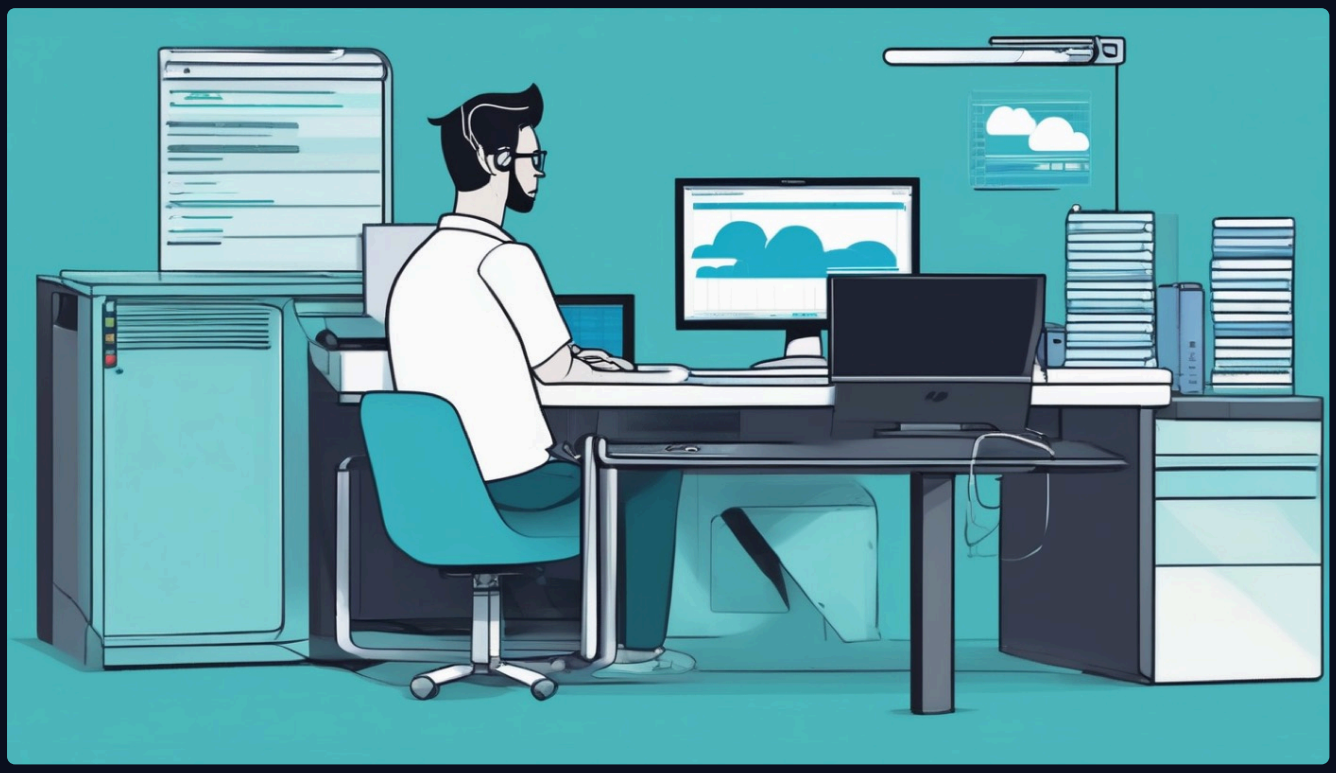
Take a worked example. A mid-size firm with no existing handling standard experiences a near-miss: a courier-delivered box of Confidential client files is left unattended in a reception area for twenty minutes. The firm writes its first handling standard afterwards, specifying that Confidential material in transit must never be left unattended and must travel in a sealed, tamper-evident container.

An exam trap here is assuming handling requirements are only about technology — encryption, access controls. Just as often the correct answer is procedural: who signs for a delivery, who is authorised to carry a device off-site, what happens if a seal is broken.

Key points: A handling standard turns a classification level into concrete storage, transport and destruction rules. The riskiest moments are when an asset leaves direct control. Procedural rules matter as much as technical ones.

Exam tip: ISC2 often tests handling with a scenario about a specific moment — a courier, a lost laptop, an unattended desk — rather than asking for the standard's contents. Match the moment to the rule that actually applies to it.

5. 2.3 Provisioning resources securely



Before any system, account or piece of software goes live, it has to be provisioned: set up, configured and handed over. That process has to leave it without weaknesses already carried into production on day one.

Secure provisioning starts with ownership: before a resource exists, someone is named as its owner, accountable for its classification and its ongoing protection. Provisioning then applies a secure baseline configuration, removes default accounts and passwords, and confirms the resource only has the access it actually needs before anyone relies on it.

This matters because once a system is live, fixing a provisioning mistake competes with every other priority. A server stood up with a default administrator password, or a cloud storage bucket created with public access left on by default, both start their working life already compromised.

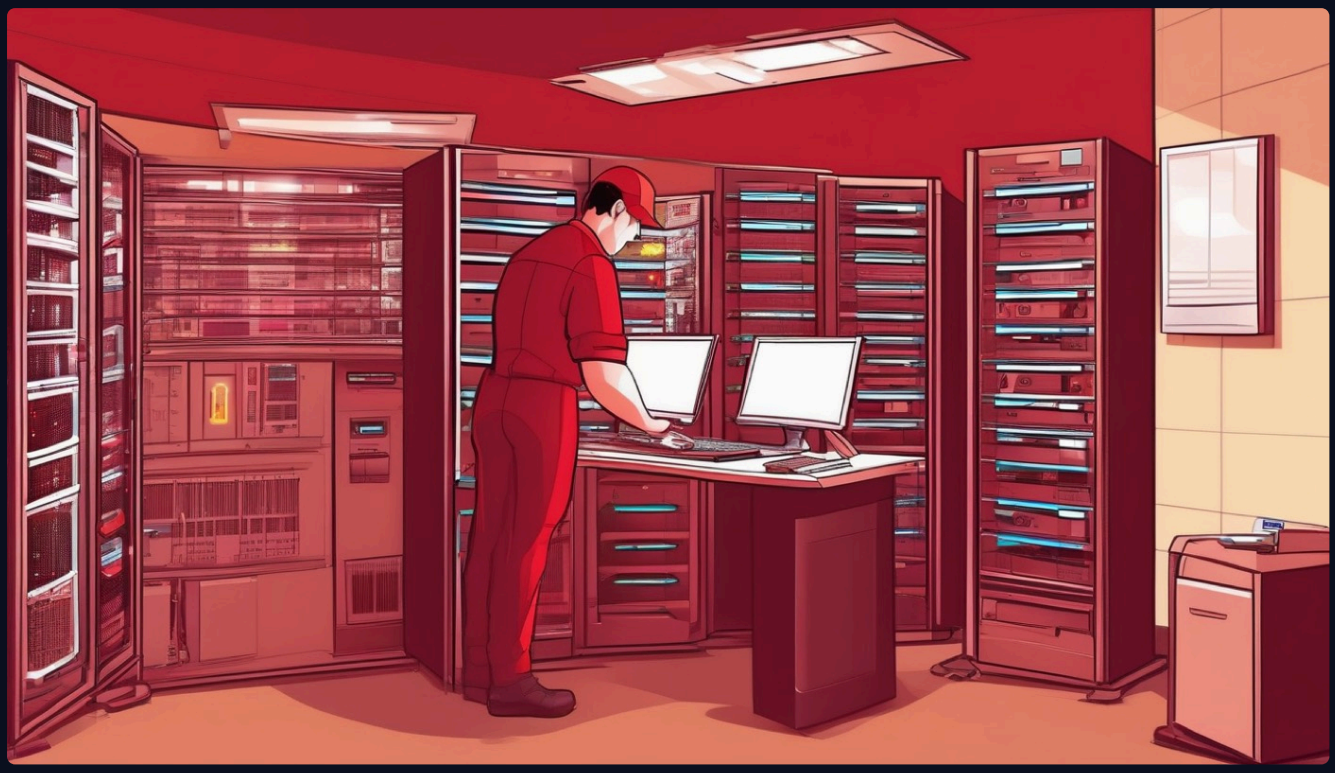
Here is how that plays out in practice. A development team onboards a new cloud server for a customer-facing application. Before go-live, the provisioning checklist assigns an owner and applies the organisation's hardened server image. It disables the default administrator account, and restricts network access to only the ports the application actually uses.

A frequent exam trap is treating "provisioning" as a one-time technical task that ends when the system boots. ISC2 treats it as a process with an owner attached. A system with no named owner has, in effect, not been provisioned securely at all, however well configured it is.

Key points: Secure provisioning assigns an owner before go-live, applies a hardened baseline, and removes default settings. The goal is a resource that starts life already meeting its classification's requirements.

Exam tip: Expect a scenario where the technical configuration looks fine but ownership is missing — that is usually the flaw ISC2 wants you to spot, because an unowned asset has nobody accountable when something goes wrong.

6. 2.3 Asset inventory and management





Provisioning a resource correctly once is not enough. An organisation also needs to know, at any later moment, exactly what it has, where it is, and whether it is still configured the way it was meant to be.

An asset inventory records every tangible asset — servers, laptops, network equipment — and every intangible asset — software licences, certificates, cryptographic keys — the organisation is responsible for. Configuration management then tracks each asset's current settings against its approved baseline, so a change can be spotted and investigated rather than discovered months later.

Licence tracking is part of the same discipline. Software is provisioned under a licence that limits how many machines or users may run it, and losing track of that limit creates both a compliance exposure and, often, an unbudgeted cost.

Walk through a worked example. An internal audit at a mid-size company finds forty installations of a licensed design application against only twenty-five purchased seats. A technician reviewing the server rack and running a laptop-based discovery scan during the audit traces the extra installs to a department that provisioned new laptops without checking the software catalogue first.

An exam trap here is assuming an inventory is complete once it is first built. ISC2 treats asset management as continuous — an inventory with no process for keeping it current is already going stale the day it is finished.

Key points: An asset inventory covers both tangible and intangible assets. Configuration management checks current state against an approved baseline. Licence tracking prevents both compliance exposure and wasted spend.

Exam tip: Expect a question where the inventory exists but is out of date — ISC2 is testing whether you know maintenance is part of the control, not whether you know what an inventory is.

7. 2.4 Data roles across the lifecycle



Every piece of data that an organisation holds has several people standing in defined relationships to it, and knowing which role applies to which person decides who is accountable when something goes wrong.

The **owner** is accountable for the data's classification and protection decisions — usually a business role, not an IT one. The **custodian** carries out day-to-day protection on the owner's instruction: backups, access control, patching. The **user** is whoever legitimately works with the data day to day. The **subject** is the person the data is about.

Data protection law adds two more precise roles that overlap with owner and custodian without being identical to them. The EU's General Data Protection Regulation defines the **controller** as whichever person or body decides why and how personal data is processed. The **processor** acts only on the controller's documented instructions.

Consider how this plays out. A company runs its own payroll system but pays an external bureau to process the monthly payslips. The company is both data owner and GDPR controller: it decides why payroll data is collected and how it is used. The bureau is the processor: it acts strictly on the company's instructions and has no right to use the data for its own purposes.

An exam trap is assuming owner and controller, or custodian and processor, are the same role under two names. They usually line up, but a company can be the controller of data a different company's system physically custodises — ownership is about accountability, custody is about who holds the keys.

Key points: Owner, custodian, user and subject are the core CBK roles. Controller and processor are GDPR's versions of owner and custodian, not always identical to them. Accountability and day-to-day protection are separate jobs.

Exam tip: ISC2 scenarios often name a vendor or a bureau and ask who is accountable. The answer is almost always the owner or controller — the party that decided why the processing happens — not whoever is physically running the system.

8. 2.4 Data lifecycle: collection



The data lifecycle begins the moment data is collected, and the decisions made at that first moment shape every protection requirement that follows.

The guiding principle is **data minimisation**: collect only what is actually needed for a stated purpose, and document that purpose at the point of collection. Data gathered just in case still has to be classified, stored, protected and eventually destroyed — it is a cost and a risk with no offsetting benefit.

Documentation matters as much as minimisation. A field collected without a recorded reason is hard to classify correctly later, and it is the first thing a privacy regulator or an internal auditor will ask about: why do you have this, and what did you tell the person it was for?

Here is one way this plays out. An audit of a company's customer sign-up form finds it asking for a date of birth and a job title, neither of which the company's marketing team can explain a use for. Both fields are removed from the form, and the remaining fields are matched, one by one, to a documented business purpose.

An exam trap is treating more data as automatically more useful, a security-neutral trade-off. ISC2 treats unexplained collection as a finding in itself, independent of whether that data is ever misused — the exposure exists the moment it is stored.

Key points: Collection should follow data minimisation — gather only what a documented purpose requires. Every collected field should trace to a stated reason. Unexplained collection is itself a risk,

not a neutral choice.

Exam tip: Expect a scenario listing several collected fields and asking which should be removed. Look for the field with no stated business purpose — that is almost always the one ISC2 wants flagged.

9. 2.4 Data lifecycle: location



Once data exists, an organisation has to know where it physically or logically sits, because location decides which country's laws, and which organisation's controls, actually apply to it.

Data residency is a choice or a requirement about where data is stored, often set voluntarily for performance or cost reasons. **Data sovereignty** is the stronger idea: a government's legal authority over data located within its borders, regardless of who owns that data or where that owner is based.

The distinction matters most when data crosses a border. A company storing customer data in a country whose laws grant broad government access to data held there faces sovereignty exposure quite separate from any technical security control it applies.

Take this scenario. A company based in one country stores its customers' personal data in a cloud region located in a second country. It then wants a team based in a third country to analyse that data. Each hop across a border has to be checked against the rules the data's country of origin imposes on transfers out – not just the destination's own rules.

An exam trap is assuming encryption solves the sovereignty problem. It does not: a government with legal authority over a location can, in many cases, compel disclosure of the data stored there regardless of how well it is encrypted, if it can also compel disclosure of the keys.

Key points: Data residency is where data is stored, often chosen for convenience. Data sovereignty is a government's legal authority over data within its borders. Cross-border transfers require checking the rules of the data's origin, not just its destination.

Exam tip: ISC2 scenarios on this topic usually involve a multinational transfer. The correct answer generally respects the origin country's transfer rules, not the destination's convenience.

10. 2.4 Data lifecycle: maintenance



Data that is collected and stored correctly still needs active maintenance. It has to be kept accurate, kept available, and recoverable if something goes wrong, or its value decays even while it sits untouched.

Backups are the core maintenance control. A **full backup** copies everything; an **incremental backup** copies only what changed since the last backup, of any type; a **differential backup** copies everything changed since the last full backup. Each trades backup speed against how long a restore takes to rebuild.

Archiving is a related but distinct idea: moving data that is no longer actively used, but still has to be retained, to cheaper, slower storage. Archived data still needs maintenance — periodic checks that it can still be read and restored — or it is not really being retained at all, just postponed as a loss.

Picture how this plays out. A company discovers, during a routine drill, that its nightly incremental backups have been running successfully for months, but the underlying full backup they depend on failed silently six weeks earlier. The incrementals are worthless without it. The fix is adding automated verification that a restore chain is actually complete, not just that each job reported success.

An exam trap is assuming a backup job reporting success means data can actually be restored. ISC2 expects you to know that only a tested restore proves a backup works — an untested backup is a belief, not a control.

Key points: Full, incremental and differential backups trade speed against restore complexity. Archiving moves inactive but retained data to cheaper storage, and still needs periodic verification. An untested backup is not a proven control.

Exam tip: Expect a scenario where backups "succeeded" but a restore still fails. The lesson ISC2 is testing is that verification, not job completion, is what proves recoverability.

11. 2.4 Data lifecycle: retention



Retention is the decision about how long data must, or may, be kept before it is destroyed, and getting that period wrong in either direction creates risk.

A retention period is set from two directions at once: a **legal minimum**, set by whichever law or regulation applies to that data, and a **business need**, how long the organisation genuinely benefits from keeping it. The correct retention period is normally the longer of the two, documented in a formal retention schedule.

Keeping data too briefly can breach a legal retention requirement or destroy evidence an organisation later needs. Keeping it too long raises cost and expands the amount of data at risk if there is a breach. For personal data, it can itself breach a privacy law that limits retention to what is necessary.

Consider this scenario. A company building its first retention schedule finds three categories of data with no agreed retention period at all: old job applicant records, expired contracts, and closed support tickets. For each, the schedule records the legal minimum that applies, the business case for keeping it longer if any exists, and a destruction date calculated from whichever is longer.

An exam trap is assuming that keeping everything forever is the safe default. ISC2 treats over-retention as a genuine control failure, not caution — every record kept past its justified need is pure downside risk with no offsetting benefit.

Key points: A retention period is the longer of the legal minimum and the genuine business need. A retention schedule documents the period and the destruction date for each data category. Over-retention is a risk, not a safe default.

Exam tip: Expect a scenario asking whether to keep data just in case. The ISC2-correct answer is almost always no, unless a specific, named legal or business reason is given.

12. 2.4 Data remanence



Data remanence is the residual representation of data that remains on storage media after an attempt has been made to remove it – and it is the reason deleting a file is not the same thing as destroying the data in it.

When a file is deleted through an ordinary operating system command, in almost every case only the pointer to that file is removed. The underlying bits stay on the media, recoverable with ordinary tools, until something else happens to overwrite that same physical location – which may never happen before the media is reused or disposed of.

This is why saying it was deleted is not a security control. The data is still physically present, in a form a reasonably capable attacker, or even consumer-grade recovery software, can retrieve. That stays true right up until the media is sanitised or destroyed by a method appropriate to its classification.

Here is a scenario worth walking through. A company sells a batch of retired laptops through a reseller after running the operating system's standard "delete" on each drive. A buyer later recovers client files, including Confidential records, using free recovery software, because deletion never actually removed the underlying data – only the pointers to it.

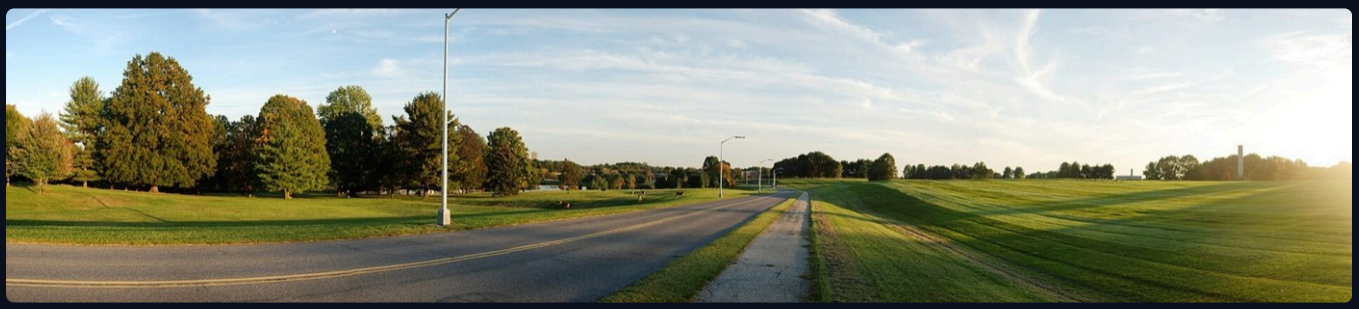
An exam trap is assuming reformatting a drive is equivalent to destroying its data. A standard format, like an ordinary delete, typically leaves recoverable remanence behind; the next chapter covers the sanitisation levels that actually address it.

Key points: Data remanence is what remains after an attempted deletion. Ordinary deletion removes only a pointer, not the underlying data. Remanence persists until a proper sanitisation or destruction method is applied.

Exam tip: ISC2 scenarios on remanence usually involve media being resold, donated or disposed of after only an ordinary delete or format. The flaw is always the same: deletion was mistaken for destruction.

13. 2.4 Data sanitisation and destruction





Properly addressing data remanence means choosing a sanitisation or destruction method that matches the classification of the data that was on the media, and the fate the media is headed for next.

NIST set out three sanitisation categories in Special Publication 800-88. **Clear** uses logical techniques, such as overwriting, to defeat ordinary, non-invasive recovery attempts. **Purge** uses stronger physical or logical techniques, such as cryptographic erasure or degaussing, to defeat recovery even with specialist laboratory equipment. **Destroy** physically destroys the media — shredding, disintegration, incineration or pulverising — so it cannot be used again at all.

This guidance comes from NIST, the US National Institute of Standards and Technology, headquartered in Gaithersburg, Maryland. Its current edition is Special Publication 800-88 Revision 2, which supersedes the 2014 original and keeps all three categories. It was finalised on 26 September 2025. What has changed is where the specific wipe techniques live. Revision 2 now points readers to

a separate standard, IEEE 2883, for choosing the right technique for a given media type, because those details change faster than NIST revises this guide.

Take a concrete scenario. An organisation is retiring a mix of media, each previously holding Confidential data: spinning hard drives, solid-state drives and a handful of magnetic backup tapes. Each needs a different method. Spinning drives can often be cryptographically purged if encrypted from new; some solid-state drives resist overwriting and may need physical destruction instead; magnetic tape responds well to degaussing, a purge-level technique.

An exam trap is assuming one method, such as degaussing, works for every kind of media. It does not. Degaussing has no effect on solid-state media at all, because there is no magnetic domain to disturb. ISC2 expects you to match method to media, not apply a favourite technique everywhere.

Key points: Clear, Purge and Destroy are the three sanitisation categories, unchanged in NIST SP 800-88 Revision 2 (2025). Technique selection for a given media type now points to IEEE 2883. Crypto-shredding is a purge-level technique that relies on destroying the encryption key, not the data itself.

Exam tip: Expect a scenario matching a media type to the right category, drawn from the NIST guidance behind every answer on this topic.

14. 2.5 Asset retention, EOL and EOS



Assets, like data, have a limited useful life, and knowing where an asset sits in that life decides whether it is still safe to keep relying on.

End-of-life (EOL) is the point at which a vendor stops selling or actively developing a product. **End-of-support (EOS)** is the later point at which the vendor stops issuing security patches altogether. Between EOL and EOS, an asset is usually still safe to run with care; after EOS, every newly discovered vulnerability in that product will simply never be fixed.

The gap between EOL and EOS can lull an organisation into a false sense of safety, because the product still works exactly as it did the day before. What has changed is invisible: the vendor's clock on patching has run out, and the risk from that point only grows as new vulnerabilities are found and never closed.

Picture a concrete scenario. A company keeps a line-of-business server running on an operating system two years after its EOL announcement and six months past its EOS date. No patches have shipped in that six months, while three new serious vulnerabilities for that operating system have been publicly disclosed. The asset's risk, measured by unpatched disclosed vulnerabilities, has risen sharply with no change to the system itself.

An exam trap is treating EOL and EOS as the same event. ISC2 expects you to know the gap between them, and to recognise that EOS, not EOL, is the date that actually matters for security risk.

Key points: End-of-life is when a vendor stops developing a product. End-of-support is when it stops patching it. Risk rises sharply after EOS, even though the asset appears unchanged.

Exam tip: Expect a scenario giving both an EOL and an EOS date and asking which matters most for security. The answer is EOS — that is the date vulnerabilities stop being fixed.

15. 2.6 Data states: at rest, in transit, in use





Data security controls are not one-size-fits-all; the right control depends on which of three states the data is in at that moment.

At rest is data sitting in storage: a server rack full of drives, a backup tape, a laptop's disk. The standard control is encryption of the storage itself. **In transit** is data moving across a network, protected chiefly by transport encryption such as TLS. **In use** is data loaded into active memory and being processed, the hardest state to protect, addressed by memory protection and increasingly by confidential-computing techniques that keep data encrypted even while a processor works on it.

Each state has a different attacker in mind. At rest, the threat is someone who gets physical or logical access to the storage itself. In transit, the threat is someone intercepting the network path. In use, the threat is someone — or something — with access to the running system itself, which is why this state is the hardest of the three to defend.

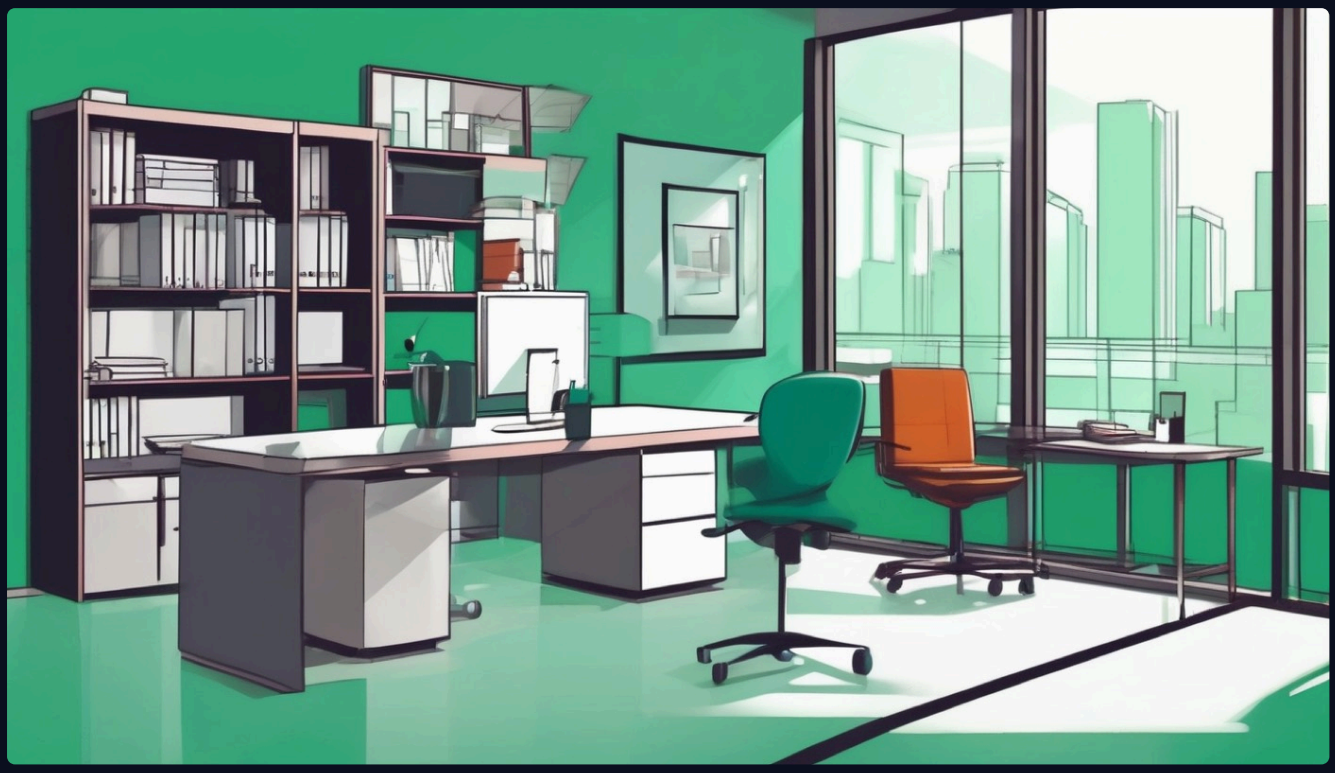
Consider a concrete case. A customer database is encrypted on the server rack where it lives at rest. It travels to an analytics application over a TLS-protected connection while in transit. It is then processed inside a confidential-computing enclave while in use, so that even an administrator with access to the running server cannot read it.

An exam trap is assuming encryption at rest alone is sufficient data protection. ISC2 expects you to name the matching control for each of the three states separately — a dataset can be perfectly encrypted at rest and completely exposed the moment it is loaded into memory for processing.

Key points: At rest, in transit and in use are the three data states. Each needs its own matching control — storage encryption, transport encryption, and memory or confidential-computing protection. The "in use" state is the hardest to defend.

Exam tip: Expect a scenario naming only one state's control and asking what is missing. ISC2 wants all three states covered, not just the most familiar one.

16. 2.6 Scoping and tailoring baselines



Security control baselines — the standard sets of controls defined by a framework — are rarely applied exactly as published; they are adjusted to fit the system and the organisation they protect.

Scoping narrows a baseline to the parts that actually apply to a given system, removing controls that address a risk the system does not have. A control for mobile device loss, for instance, gets scoped out of a baseline applied to a server with no mobile component at all. **Tailoring** adjusts the controls that remain, setting parameters or substituting a compensating control where the standard one cannot be applied as written.

The two operations are sequential and distinct. Scoping decides which controls belong in the baseline for this system at all; tailoring then adjusts how the controls that survived scoping are actually implemented.

Here is how this looks in practice. A company applies its main security baseline to a small branch office with no dedicated IT staff on site. Scoping removes controls assuming a local system administrator is present. Tailoring then substitutes a compensating control for the controls that remain but cannot be met exactly — centralised remote monitoring in place of a control written around local log review.

An exam trap is confusing tailoring with simply skipping a control that is inconvenient. Tailoring requires a documented, equivalent compensating control; dropping a control with nothing in its place is not tailoring, it is an unmanaged gap.

Key points: Scoping removes controls that do not apply to a given system. Tailoring adjusts the controls that remain, often through a documented compensating control. Dropping a control with no substitute is a gap, not tailoring.

Exam tip: ISC2 scenarios usually test whether you can tell scoping from tailoring. Scoping answers whether a control applies at all. Tailoring answers how to actually implement the ones that do.

17. 2.6 Selecting data security standards



Choosing which security standard applies to a given dataset is driven by the industry the organisation operates in and the type of data it holds, not by a single, universal checklist.

A retailer or any organisation handling card payment data must map its controls to the Payment Card Industry Data Security Standard, now at version 4.0.1. Every one of its requirements, including those once given a future compliance date, has been mandatory since 31 March 2025. A healthcare organisation maps instead to health-data-specific regulation. A government contractor maps to the standards its contract specifies.

The selection decision usually starts from the data itself, not the organisation's size or sector label: does the dataset include card numbers, health records, or government-classified information? That answer points directly to the standard that has to apply, often alongside, not instead of, a general framework the organisation already follows.

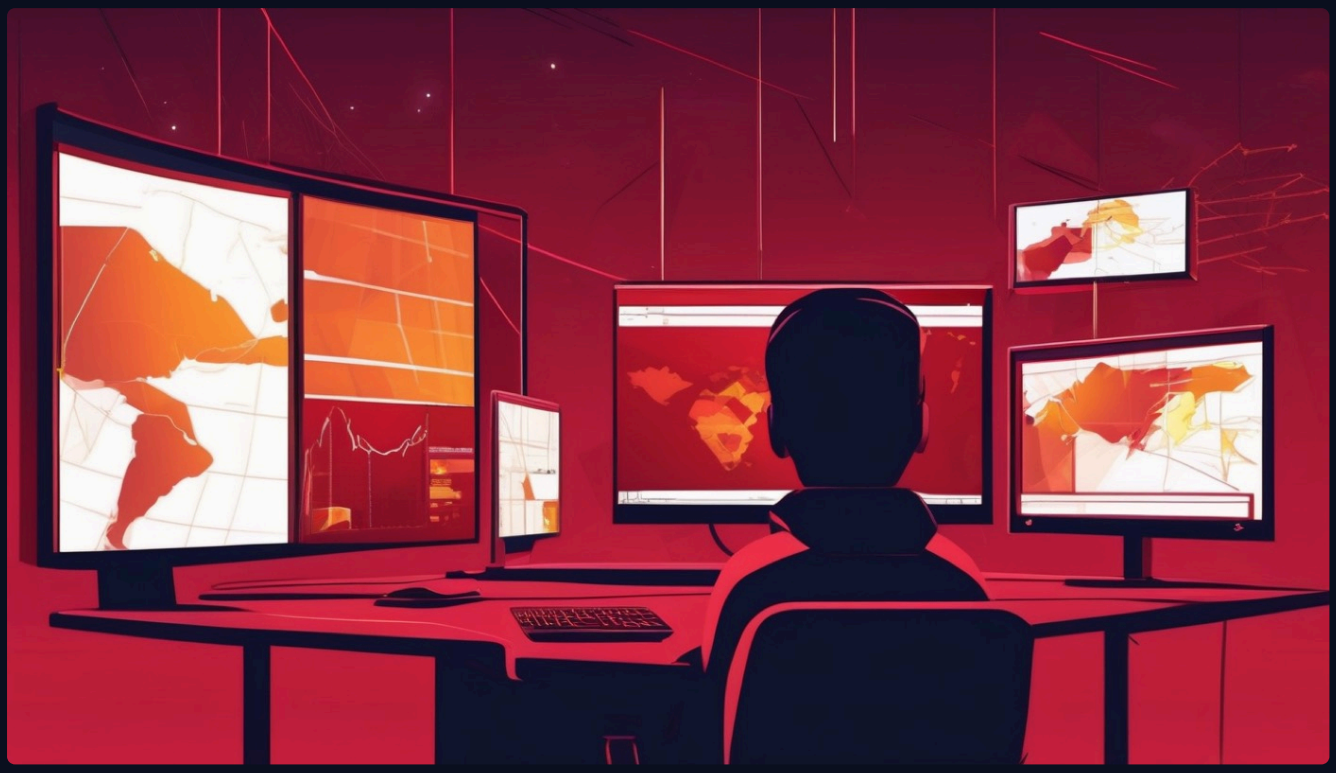
Walk through this scenario. An online retailer that already follows a general security framework for its internal systems adds card payment processing to its website. The card data triggers a separate, mandatory mapping to PCI DSS v4.0.1 for every system that stores, processes or transmits that data. The general framework alone does not satisfy that specific, binding requirement.

An exam trap is assuming one strong general framework is "enough" regardless of the data involved. ISC2 expects you to recognise that specific data types trigger specific, non-negotiable standards layered on top of whatever general framework is already in place.

Key points: Standard selection follows the data, not just the sector. PCI DSS v4.0.1 is mandatory in full since 31 March 2025 for card payment data. A general framework does not substitute for a data-specific mandatory standard.

Exam tip: Expect a scenario naming a specific data type and asking which standard applies. Match the data to its mandatory standard first, then layer the organisation's general framework around it.

18. 2.6 Data protection: DRM, DLP, CASB



Beyond encryption and access control, three further tools focus specifically on controlling what happens to data as it moves and as it is used.

Digital Rights Management (DRM) controls what someone may do with content even after it has left the owner's direct systems. It can restrict copying, printing or forwarding of a protected document, for instance. **Data Loss Prevention (DLP)** watches data in use and in transit for patterns that match sensitive information, and blocks or flags an unauthorised attempt to move it out of the organisation. A **Cloud Access Security Broker (CASB)** sits between an organisation's users and the cloud services they use. It enforces policy, such as blocking an unapproved file-sharing service, that the cloud provider itself does not enforce.

Each tool addresses a different point in the data's journey. DRM travels with the content itself, wherever it goes. DLP watches the organisation's own boundary. CASB watches the gateway to cloud services specifically, which is where an organisation's own DLP tooling often has the least visibility.

Take this worked scenario. An employee attempts to upload a spreadsheet of Confidential customer records to a personal cloud storage account from a work laptop. The organisation's DLP tooling recognises the pattern of data in the file and blocks the upload. Its CASB, watching cloud traffic separately, flags the same attempt as use of an unapproved cloud service. Security gets two independent signals of the same incident.

An exam trap is treating DRM, DLP and CASB as interchangeable data protection tools. ISC2 expects you to know which one applies to which problem. DRM is for content leaving the organisation's

control entirely. DLP is for movement across the organisation's own boundary. CASB is specifically for cloud service use.

Key points: DRM protects content after it leaves the owner's systems. DLP watches for unauthorised movement of sensitive data. A CASB enforces policy specifically at the boundary with cloud services.

Exam tip: Expect a scenario describing an incident and asking which tool would have caught it. Match the incident's location – inside the content, at the organisation's boundary, or at the cloud gateway – to the matching tool.

19. 2.6 Protection methods: tokenisation, masking, anonymisation



Sometimes data has to be used for a purpose that does not need the real value: testing, analytics, display to a partially authorised user. Several techniques reduce exposure in those cases without destroying the data's usefulness.

Tokenisation replaces a sensitive value with a non-sensitive substitute, or token. The mapping between the two is kept in a separate, tightly controlled vault, and the original value can be retrieved only through that vault. **Masking** hides part of a value for display – showing only the last four digits of a card number – while the full value remains stored elsewhere unchanged. **Pseudonymisation** replaces identifying fields with a key that can be reversed by someone holding the right additional information. **Anonymisation** removes identifying detail so thoroughly that it cannot be reversed at all, by anyone.

The distinction that matters most for an exam, and for real compliance decisions, is reversibility. Tokenisation and pseudonymisation are reversible, by someone holding the key or the vault. Anonymisation, done properly, is not reversible by anyone. That is also why genuinely anonymised data falls outside the scope of most privacy law, while pseudonymised data does not.

Picture this scenario. A company needs a realistic copy of its customer database for a software testing team that has no business need to see real customer details. Names and addresses are tokenised, so testers see consistent but fake-looking values. Card numbers are masked to their last four digits. A

separate analytics extract, stripped of every identifying field with no key retained anywhere, is treated as anonymised and handled with lighter controls as a result.

An exam trap is treating pseudonymised data as if it were anonymised, and therefore outside privacy law's scope. It is not — if a key exists anywhere that could reverse it, the data is still personal data under most privacy law, however well it has been obscured for everyday use.

Key points: Tokenisation and pseudonymisation are reversible, by someone holding the key. Masking hides part of a value for display while keeping the full value elsewhere. Anonymisation, done properly, is irreversible and changes a dataset's legal status.

Exam tip: Expect a question asking whether a given technique still counts as personal data under privacy law. If a key exists anywhere to reverse it, the answer is yes.

20. 2.6 Privacy obligations across the lifecycle





Privacy law does not sit in one place in the data lifecycle. It runs through classification, handling, retention and destruction all at once. Domain 2's controls are, in large part, how an organisation actually meets it.

The European Union's General Data Protection Regulation, enforceable since 25 May 2018 and enforced across the EU by national authorities under the oversight of the European Commission in Brussels, is the privacy law most referenced across this domain. It requires personal data to be classified and protected according to its sensitivity, collected only for a stated purpose, retained no longer than necessary, and properly destroyed or anonymised at the end of its justified life.

Several of GDPR's rights depend directly on the asset security practices already covered in this module. Take a request from an individual to know, correct or erase their own data. An organisation cannot answer it quickly, or at all, without already knowing where that person's data lives, who the custodian is, and what its retention status is.

Consider this worked scenario. A customer submits a data subject access request, asking a company to disclose everything it holds about them. The company's data lifecycle practices already tie every dataset to an owner, a location and a retention schedule. So it compiles and verifies the full disclosure within the statutory deadline – something a company with no inventory or lifecycle tracking at all would struggle to do.

An exam trap is treating privacy as a separate legal function from asset security. ISC2 expects you to see classification, handling, retention and destruction as the practical machinery that makes privacy compliance possible, not as a parallel track run by a different department.

Key points: GDPR's requirements map directly onto classification, collection, retention and destruction. Individual rights such as access and erasure depend on practices already covered in this module. Privacy compliance and asset security are the same work, not two separate ones.

Exam tip: Expect a scenario about responding to a privacy request, where the real test is whether the organisation's asset security practices make that response possible at all.

21. From the old 10-domain CBK to today's Domain 2



Domain 2 did not appear from nothing. ISC2 restructured the CISSP Common Body of Knowledge from ten domains down to eight in April 2015, and today's Asset Security domain absorbed material from two of the older domains.

Parts of the old *Information Security Governance and Risk Management* domain — particularly its treatment of data classification and asset ownership — moved into what is now Domain 2. Parts of the old *Operations Security* domain moved into Domain 2 as well, particularly media handling, data remanence and secure disposal, alongside material that went to other domains entirely.

Knowing this history is not just trivia. It explains why Domain 2 reads, in places, like two different topics stitched together. Classification and ownership on one side, and sanitisation and destruction on the other, are exactly the two older domains Domain 2 inherited, carried forward into the 2024 outline that is examined today.

Candidates who studied older material, or whose employer trained them years ago under the ten-domain structure, sometimes still think of data remanence as an "operations" topic rather than an "asset security" one. The content has not changed much; its address in the CBK's structure has.

Key points: The ten-domain CBK became eight domains in April 2015. Today's Domain 2 absorbed classification and ownership material from the old governance domain and sanitisation and disposal material from the old operations domain.

Exam tip: ISC2 does not test the 2015 history directly, but recognising that Domain 2 spans two older disciplines explains why its sub-domains feel less tightly unified than some other domains — and helps you place a topic correctly when a question's framing feels unfamiliar.

22. Domain 2 full recap



Here is the whole of Domain 2, tied back together.

Everything starts with classification: sorting information and assets by the harm their loss would cause, using government or commercial labels, then labelling and marking them so that classification is visible to whoever handles them next. An asset register and configuration management keep track of what exists and whether it still matches its approved state.

Four roles — owner, custodian, user and subject — define who is accountable, who protects day to day, who uses the data, and who the data is about. GDPR's controller and processor are the privacy-law versions of owner and custodian.

The data lifecycle then runs from collection, through tracking its location, maintaining and backing it up, and retaining it for exactly as long as it is needed. It ends in sanitising or destroying it, because deletion alone leaves recoverable remanence behind. Only Clear, Purge or Destroy, chosen to match the media, actually removes it.

Assets, not just data, have a lifespan too. End-of-life and end-of-support mark the points where a vendor stops developing, then stops patching, a product, and risk rises sharply once support ends. Controls then have to match whichever of the three data states applies at that moment: at rest, in transit, in use. They are scoped and tailored to the system they protect, and mapped to whatever mandatory standard the data itself triggers.

DRM, DLP and CASB each guard a different point in data's journey outward. Tokenisation, masking, pseudonymisation and anonymisation each trade off usefulness against reversibility. Privacy law, chiefly GDPR, runs through every one of those controls at once, which is why Domain 2's practices are what actually make privacy compliance possible.

That is the whole domain: ten per cent of the exam, and the foundation that classification-driven decisions across every other domain rest on.

23. Review questions 1: classification, roles and handling



Time to test what you know. Each question gives four options, a short pause to think, then the answer and why the others are wrong.

Question one. A company rates its payroll records as its highest commercial classification tier. Which factor most directly justified that choice? A: the file size. B: the number of years the records are kept. C: the harm disclosure would cause to employees. D: the department that created the file.

The answer is C. Classification is driven by the harm that loss or disclosure would cause, not by size, retention length or which department created the data. A, B and D each describe the record without touching the harm a breach would cause.

Question two. An unlabelled laptop is found holding Confidential data with no entry in the asset register. What is the most complete fix? A: add a classification sticker only. B: delete the data to remove the risk. C: add the laptop to the register with an owner and a review date, and label it. D: encrypt the laptop and take no further action.

The answer is C. A sticker alone, as in A, fixes visibility but not accountability. Deleting the data, as in B, may breach retention obligations. Encrypting, as in D, is good practice but leaves the asset untracked. Only C closes both the labelling gap and the register gap together.

Question three. Under GDPR, a company decides why customer data is collected and how it will be used, while an external bureau processes that data strictly on the company's instructions. What is the bureau's role? A: controller. B: processor. C: subject. D: custodian only, with no legal role.

The answer is B. The company deciding purposes and means is the controller, and the bureau acting on instructions is the processor. A describes the company, and C is the person the data is about. D is wrong because GDPR gives the processor a defined legal role, not just a technical one.

Question four. A handling standard for Confidential material requires it never be left unattended in transit. Which incident would this rule most directly have prevented? A: a server running an outdated patch. B: a courier-delivered box left unattended in reception. C: a licence count exceeding purchased seats. D: a backup job failing silently.

The answer is B. Patching, licensing and backup failures, in A, C and D, are real problems but belong to other sub-domains entirely. Only B describes an asset left unattended in transit, exactly the scenario the rule addresses.

Question five. A security manager is told a new cloud server is fully configured and ready to go. What question best tests whether it was securely provisioned? A: how much storage does it have? B: who is the named owner, and were default accounts removed? C: what is its IP address? D: how fast is its network connection?

The answer is B. Storage, IP address and network speed, in A, C and D, are configuration details unrelated to security provisioning. B tests the two things that actually define secure provisioning: accountable ownership and removal of default weaknesses.

Question six. An internal audit finds forty software installations against twenty-five purchased licence seats. What practice failure does this most directly point to? A: data classification. B: licence tracking within asset management. C: data residency. D: data remanence.

The answer is B. Classification, residency and remanence, in A, C and D, concern different problems entirely; the mismatch between installs and purchased seats is a licence-tracking failure, part of ongoing asset management.

That is review set one. On to the data lifecycle.

24. Review questions 2: lifecycle, retention and destruction



Question seven. A sign-up form collects a date of birth with no documented business reason. What does ISC2 expect you to flag? A: nothing, since the field is harmless. B: the unexplained collection itself, independent of any later misuse. C: only the storage cost. D: only whether the field is encrypted.

The answer is B. Encryption and storage cost, in C and D, are secondary; A ignores that unexplained collection is itself a finding. The exposure exists the moment unneeded data is collected, whether or not it is later misused.

Question eight. A company transfers customer data across three countries for processing. Whose rules govern each cross-border hop? A: only the destination country's rules. B: only the company's home country's rules. C: the data's country of origin's transfer rules, checked at each hop. D: whichever country has the weakest rules.

The answer is C. A, B and D each pick one jurisdiction regardless of the data's own origin; the correct approach checks the origin country's transfer rules at every border the data crosses, not just the final destination.

Question nine. Nightly incremental backups have run successfully for months, but the full backup they depend on failed silently six weeks ago. What does this reveal? A: incremental backups are inherently unreliable. B: a successful job report does not prove a working restore chain. C: differential backups should have been used instead. D: the backup schedule was too frequent.

The answer is B. Incremental backups are not inherently unreliable, which rules out A. Differential backups and scheduling frequency, in C and D, do not address the actual fault. Only a tested restore proves the chain works — a job's success report alone does not.

Question ten. A retention schedule sets a destruction date for a data category as the longer of two figures. What are those two figures? A: file size and storage cost. B: the legal minimum and the genuine business need. C: the classification level and the owner's preference. D: the backup frequency and the archive location.

The answer is B. File size, classification level and backup frequency, in A, C and D, play no role in setting a retention period. It is set from the longer of the legal minimum retention and the organisation's genuine business need.

Question eleven. A retired laptop's drive was "deleted" using the operating system's standard command before resale. A buyer later recovers client files. What actually happened? A: the deletion command failed to run. B: only the file pointers were removed; the underlying data remained. C: the recovery software exploited a vulnerability. D: the files were never really deleted data at all.

The answer is B. The deletion command ran correctly, ruling out A. No vulnerability was needed, ruling out C. The recovered files were genuinely deleted data, ruling out D. Deletion only removes the pointer, leaving the underlying bits recoverable.

Question twelve. An organisation needs to sanitise a mix of spinning hard drives, solid-state drives and magnetic tape, all previously holding Confidential data. Why can't one method be used for all three? A: cost differs by media type. B: each media type responds differently to a given technique — degaussing, for instance, has no effect on solid-state media. C: only tape requires sanitisation at all. D: NIST SP 800-88 only covers hard drives.

The answer is B. Cost, in A, is a real factor but not the reason one method fails. C is wrong because every media type here needs sanitisation. D is wrong because NIST's guidance covers all media types. Degaussing works through magnetic fields, which solid-state media simply does not have.

That completes lifecycle, retention and destruction. The final set covers controls, privacy and how Domain 2 connects to the rest of the exam.

25. Review questions 3: controls, privacy, other domains



Question thirteen. A customer database is encrypted at rest, sent over TLS in transit, but processed in plain, unencrypted memory while in use. What is the gap? A: at-rest protection is missing. B: in-transit protection is missing. C: in-use protection is missing, the hardest state to defend. D: there is no gap; two out of three states is sufficient.

The answer is C. At rest and in transit are both correctly protected here, ruling out A and B. D is wrong because all three states need matching controls. The gap is in-use protection, which needs memory protection or confidential-computing techniques.

Question fourteen. A branch office removes controls assuming on-site IT staff, then substitutes centralised remote monitoring for the controls that remain. What two operations does this describe, in order? A: tailoring, then scoping. B: scoping, then tailoring. C: classification, then labelling. D: provisioning, then retention.

The answer is B. Classification and labelling, and provisioning and retention, in C and D, are unrelated operations. Scoping removes controls that do not apply, and happens first. Tailoring then adjusts what remains, and happens second — the reverse order in A is wrong.

Question fifteen. A retailer adds card payment processing to its website, already covered by a general security framework. What additional, mandatory standard applies? A: none; the general framework is sufficient. B: PCI DSS v4.0.1, specifically for systems handling card data. C: GDPR only, since card data is personal data. D: whichever standard is cheapest to implement.

The answer is B. A general framework, as in A, does not satisfy a mandatory, data-specific standard. GDPR, in C, may also apply but does not replace PCI DSS. Cost, in D, is not the selection criterion. Card data specifically triggers PCI DSS v4.0.1.

Question sixteen. An employee's spreadsheet upload to personal cloud storage is blocked by data loss prevention tooling and separately flagged by a cloud access security broker. Why does the organisation get two signals? A: it is a wasteful duplication with no value. B: DLP watches the organisation's boundary, while the CASB watches the cloud service gateway specifically. C: both tools do exactly the same job. D: only the CASB signal is reliable.

The answer is B. The two signals are not duplication, ruling out A. The tools do not do the same job, ruling out C. Both signals are reliable, ruling out D. DLP and CASB watch different points in the data's path, which is why one incident can trigger both independently.

Question seventeen. A pseudonymised dataset still has a key held in a separate system that could reverse it back to identifiable data. Under most privacy law, how is this dataset treated? A: anonymised, and outside the law's scope. B: still personal data, because it remains reversible. C: neither anonymised nor personal data. D: exempt, because the key is stored separately.

The answer is B. A, C and D all assume the separation or storage of the key changes the data's legal status — it does not. If a key exists anywhere that could reverse it, the data remains personal data, whatever technique was applied to obscure it.

Domain 2 connects directly to the domains either side of it. Domain 1's governance and risk management sets the policies that classification and retention schedules implement. Domain 3's security architecture decisions depend on knowing which data states and classifications a system must protect. Domain 7's operations security runs the backup, monitoring and sanitisation processes that Domain 2 defines the rules for. A security professional who has not mastered asset security is applying every other domain's controls without knowing what, exactly, they are protecting.

That is the whole of Domain 2: ten per cent of the exam, and the foundation the other seven domains build on.

Sources and credits

Primary sources

- ISC2, 'CISSP Certification Exam Outline', effective 15 April 2024 (isc2.org) - Domain 2 weighted 10%; the six sub-domain structure (2.1-2.6). Re-checked by web search 11 October 2026 (isc2.org's own page, corroborated by Readynez, Balanced Security and InfoSecTrain's 2024/2026 guides).
- ISC2's own published history of the CISSP CBK restructure from ten domains to eight domains, April 2015 - today's Domain 2 absorbed classification/ownership material from the former Information Security Governance and Risk Management domain and media-handling/remanence/disposal material from the former Operations Security domain.
- NIST Special Publication 800-88 Revision 1, 'Guidelines for Media Sanitization', finalised December 2014 - defined the Clear, Purge and Destroy categories.
- NIST Special Publication 800-88 Revision 2, 'Guidelines for Media Sanitization', finalised 26 September 2025 (csrc.nist.gov/pubs/sp/800/88/r2/final) - supersedes Revision 1, keeps the three categories, and redirects technique selection to IEEE 2883.; this is the session's one 'latest news' correction to the originally scoped facts.
- Regulation (EU) 2016/679 (General Data Protection Regulation), enforceable from 25 May 2018 - the controller and processor definitions. Re-checked by web search 11 October 2026 (Wiley, Wolters Kluwer, Lexgo summaries agree on the date and definitions).
- PCI Security Standards Council, PCI DSS v4.0.1, published 11 June 2024, replacing v4.0 on 31 December 2024. Re-checked by web search 11 October 2026 (usd.de, Secureframe, KPMG summaries): all requirements, including the 51 once future-dated as best practice, became mandatory on 31 March 2025.
- ISC2, 'CISSP Certification Exam Outline' (2024) - exam format: Computerized Adaptive Testing, English exam, 100-150 questions, 3-hour limit, pass mark 700/1000.

Photo credits (Wikimedia Commons)

- NIST headquarters, Gaithersburg, Maryland (2008) - the agency behind SP 800-88 media sanitisation guidance: Owenusa, Public domain - https://commons.wikimedia.org/wiki/File%3ANist_2.jpg
- Berlaymont building, Brussels - headquarters of the European Commission, which enforces GDPR: dimitrisvetsikas1969, CC0 - https://commons.wikimedia.org/wiki/File%3AEuropean_Commission.jpg
- Server rack close-up, NERSC data center: Derrick Coetzee from Berkeley, CA, USA, CC0 - https://commons.wikimedia.org/wiki/File%3ARear_of_rack_at_NERSC_data_center_-_closeup.jpg
- Technician working on a server rack with a laptop, NERSC data center: Derrick Coetzee from Berkeley, CA, USA, CC0 - https://commons.wikimedia.org/wiki/File%3ATechnician_with_laptop_working_on_server_rack_at_NERSC.jpg
- National Operations Center, US Department of Homeland Security: Unknown authorUnknown author, Public domain - https://commons.wikimedia.org/wiki/File%3ANOCNationalOperationsCenterHomelandSecurityUSA.jpg

Not regulated financial advice.