GDPRGARD.EU · FREE GUIDE

GDPR Cybersecurity *Guide.*

Technical and organisational measures for Article 32 compliance — encryption, access control, breach detection, the 72-hour notification rule, vendor security, and a practical 20-point checklist.

GDPRGard.eu

2026 Edition

Table of Contents

1. Why Cybersecurity Is a Legal Obligation, Not Just an IT Concern

Most businesses treat cybersecurity and data protection as two separate departments — one runs firewalls and backups, the other reviews contracts and writes privacy policies. Under the GDPR, they are the same obligation.

Article 5(1)(f) of the GDPR requires that personal data be processed in a manner that ensures "appropriate security," including protection against unauthorised or unlawful processing and against accidental loss, destruction, or damage. Article 32 then turns that principle into a concrete legal duty: implement technical and organisational measures appropriate to the risk.

This matters because GDPR fines for security failures are not rare edge cases. A ransomware attack, a misconfigured cloud storage bucket, or a phishing email that succeeds because staff were never trained are not just IT incidents — they are compliance failures with the same enforcement exposure as an unlawful marketing list or a botched data subject access request. Regulators have fined organisations under Article 32 specifically for inadequate encryption, weak access controls, and delayed patching, independent of whether any other GDPR article was breached.

For a small or mid-sized business, this guide's purpose is narrow and practical: to translate Article 32's legal language into the specific measures a real organisation can implement, in order of priority, without needing an in-house security team.

> **Who this guide is for:** businesses with an internal IT function or an external IT provider, but no dedicated security or compliance department — the profile of most GDPRGard readers. If you already have a mature information security management system (ISO 27001, SOC 2, or equivalent), most of this will already be in place; use it as a cross-check against Article 32 specifically.

2. Article 32: What "Appropriate" Security Actually Means

Article 32(1) GDPR requires controllers and processors to implement technical and organisational measures to ensure a level of security "appropriate to the risk," and lists four examples without making any of them mandatory in every case:

1. **Pseudonymisation and encryption** of personal data
2. **Ongoing confidentiality, integrity, availability and resilience** of processing systems and services
3. **The ability to restore availability and access** to personal data in a timely manner in the event of a physical or technical incident
4. **A process for regularly testing, assessing and evaluating** the effectiveness of technical and organisational measures

The word "appropriate" is doing real legal work here, and it is where most businesses go wrong in both directions. Article 32(1) is explicit that appropriateness must weigh:

- **The state of the art** — what security measures are realistically available and commonly used today, not what was standard five years ago.
- **The cost of implementation** — a five-person startup is not held to the same standard as a bank, but "we couldn't afford it" is not a defence if the cost of basic measures (multi-factor authentication, for instance) is genuinely low.
- **The nature, scope, context and purposes of processing** — a CRM holding contact names looks different from a platform processing health records or children's data.
- **The risk to individuals** — likelihood and severity of harm if the data were exposed, altered, or lost.

In practice, this means there is no single checklist that satisfies Article 32 for every organisation — the same way there is no single insurance policy that fits every business. But there is a well-established floor of measures that regulators, auditors, and courts have repeatedly treated as the baseline "appropriate" standard for any business handling ordinary personal data over the internet. The rest of this guide covers that floor, plus what changes as your risk profile rises.

Article 32 also applies to processors, not just controllers. If you provide software, hosting, or services that process personal data on behalf of clients, your own security posture is directly regulated — independent of what your client's privacy policy says.

3. Encryption and Pseudonymisation

Encryption is the single most frequently cited Article 32 measure, and for good reason: it is the one control that keeps working even after every other control has failed. If a laptop is stolen, a backup tape is lost, or a database is exfiltrated, encryption is what stands between "we had an incident" and "we had a reportable breach."

In transit. Every connection carrying personal data — website forms, API calls, admin panels, database connections between servers — should run over TLS 1.2 or later. This is no longer a security "nice to have"; browsers actively flag unencrypted (HTTP) connections to users, and Let's Encrypt makes certificates free. There is no remaining cost justification for skipping this.

At rest. Databases, file storage, and backups containing personal data should be encrypted at the storage layer (full-disk or volume encryption at minimum), with database-level or field-level encryption added for especially sensitive fields — passwords (always hashed, never encrypted — see below), payment details, health data, or national ID numbers. Most managed cloud databases (AWS RDS, Azure SQL, Google Cloud SQL) offer this as a checkbox at provisioning time, at negligible cost.

On endpoints. Laptops and mobile devices that can access personal data should have full-disk encryption enabled — BitLocker on Windows, FileVault on macOS, both built in and free. A stolen laptop with an unencrypted disk is a reportable breach; a stolen laptop with disk encryption enabled generally is not, because the data was rendered unintelligible to whoever took it.

Password storage is a special case. Passwords must never be encrypted (encryption is reversible with the right key) — they must be hashed with a slow, purpose-built algorithm (bcrypt, scrypt, or Argon2) with a unique salt per password. If your system can email a user their original password, or if a breach exposed passwords in plaintext or in a fast general-purpose hash (MD5, SHA-1), that is itself a serious Article 32 finding, independent of how the breach occurred.

Pseudonymisation is a distinct, complementary technique: replacing identifying fields with a token or reference number, so that the data cannot be attributed to a specific person without additional information held separately. It does not make data anonymous (pseudonymised data is still personal data under the GDPR), but it meaningfully reduces risk — a leaked analytics dataset keyed by customer ID rather than name and email is a materially smaller incident than the alternative, and Article 32 explicitly rewards this as a risk-reduction measure.

4. Access Control and Identity Management

Most real-world breaches do not start with a sophisticated exploit — they start with a valid login that should not have worked, or should not have had access to what it reached. Access control is where Article 32's "confidentiality" requirement lives day to day.

The principle of least privilege. Every account should have the minimum access needed to do its job, and no more. In practice this means: not every employee needs admin rights to the CRM; not every developer needs production database access; a former employee's accounts should be disabled the day they leave, not "at some point." Audit who has access to systems holding personal data at least twice a year, and remove anything that is not justified by a current role.

Multi-factor authentication (MFA). This is the single highest-value, lowest-cost security control available today. Credential-stuffing and phishing attacks succeed against passwords alone constantly; they very rarely succeed against an account also protected by an authenticator app or hardware key. MFA should be mandatory for: any admin or superuser account, any account with access to customer personal data, any remote access (VPN, cloud console) and any email account (email is usually the master key to password resets for everything else). Nearly every SaaS platform now offers MFA for free — the remaining barrier is almost always habit, not cost, and regulators increasingly treat its absence on privileged accounts as a straightforward Article 32 failure.

Role-based access control (RBAC). Rather than granting access person by person, define roles (support agent, finance, engineering, admin) with a fixed permission set, and assign people to roles. This makes access reviews tractable — you audit ten roles instead of two hundred individual grants — and makes over-privileging visible.

Shared accounts and shared credentials (a single "admin" login used by the whole support team, a password shared over Slack) should be eliminated wherever the underlying system supports individual accounts. Shared credentials break auditability — if something goes wrong, you cannot determine who did it — and they cannot be individually revoked when one person leaves.

Session management. Admin sessions should time out after a period of inactivity, and systems handling sensitive data should log and, ideally, alert on unusual access patterns — a support account suddenly exporting the full customer database at 3 a.m. is exactly the kind of event access logging exists to catch.

5. Security by Design and by Default (Article 25)

Article 25 requires that data protection be built into systems and processes from the outset, not bolted on afterward — and it has a direct security dimension that is easy to overlook.

By design: when building or commissioning a new system that will process personal data — a new signup form, a new internal tool, a new integration with a third-party API — security requirements (encryption, access control, data minimisation) should be part of the initial specification, not a post-launch remediation item. Retrofitting security into a live system is always more expensive and more error-prone than designing it in.

By default: systems should default to the most privacy- and security-protective configuration without requiring the user to actively choose it. A new user's profile should not default to "public"; a new integration should not default to requesting broader data access than it needs; a new cloud storage bucket should not default to public read access (this specific default has caused a large share of the "misconfigured S3 bucket" breach headlines of the last decade — the fix is almost always a one-line configuration change, applied too late).

Data minimisation as a security control. The GDPR's data minimisation principle (Article 5(1)(c)) is also a practical security measure: data you never collected cannot be breached. Before adding a new field to a form or a new data export to a report, ask whether it is actually needed — every unnecessary field is unnecessary risk sitting in your database.

6. Detecting a Breach: What Counts and What Doesn't

A "personal data breach" under Article 4(12) GDPR is broader than most people assume: it is any breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data. This covers far more than a hacker stealing a database. It also includes:

- An email containing customer data sent to the wrong recipient
- A laptop or phone with unencrypted data being lost or stolen
- A ransomware attack that encrypts (denies availability to) personal data, even if nothing is exfiltrated
- An employee accessing records outside their job function out of curiosity
- A misconfigured system exposing data publicly, even briefly, even if you cannot prove anyone accessed it
- A software bug that shows one customer another customer's data

You cannot report what you don't detect, so detection capability is itself an Article 32 obligation (the "ongoing… resilience" and "regularly testing" requirements). At minimum, this means: centralised logging for systems handling personal data, retained long enough to investigate an incident (90 days is a common baseline); alerting on anomalous access or export volumes; and a defined internal process for staff to report a suspected incident immediately, without fear of blame — the single biggest driver of late or missed breach notifications is an employee who noticed something wrong and didn't know who to tell, or was afraid to.

7. The 72-Hour Rule: Breach Notification Under Articles 33 and 34

Once a controller becomes aware of a personal data breach, Article 33 requires notification to the relevant supervisory authority **without undue delay and, where feasible, not later than 72 hours** after becoming aware of it — unless the breach is unlikely to result in a risk to individuals' rights and freedoms.

"Becoming aware" starts the clock, not "confirming the full scope." Regulators have been consistent on this: you do not get to delay notification until your investigation is complete. If you know enough to conclude a breach has probably occurred, the clock is running, and Article 33(4) explicitly allows notification "in phases" — report what you know within 72 hours, and supplement it as the investigation continues.

What the notification to the supervisory authority must include (Article 33(3)):

- The nature of the breach, including categories and approximate number of data subjects and records affected
- The name and contact details of your data protection officer or other contact point
- The likely consequences of the breach
- The measures taken or proposed to address the breach, including, where appropriate, measures to mitigate its possible adverse effects

If you miss 72 hours, you are not automatically non-compliant — but Article 33(1) requires that a late notification be accompanied by reasons for the delay, and "we were still investigating" is a weaker justification than a genuine cause outside your control. This is precisely why a pre-built incident response plan (Section 8) matters: organisations that have to design their notification process for the first time during an actual breach routinely miss the deadline.

When individuals must also be told (Article 34). Beyond notifying the regulator, you must also communicate the breach directly to affected individuals, without undue delay, if the breach is likely to result in a **high risk** to their rights and freedoms — for example, exposed passwords, financial data, health data, or data that enables identity theft. This notification must be in clear, plain language and cover the same substantive points as the regulator notification. Article 34(3) provides limited exceptions — notably, if the exposed data was encrypted with a key that was not also compromised (which is exactly why Section 3 matters so much: strong encryption can be the difference between a regulator-only notification and having to email every affected customer).

Processors have a parallel, faster duty. Under Article 33(2), a processor must notify the controller "without undue delay" after becoming aware of a breach — the controller's 72-hour clock runs from when either the controller or its processor became aware, so a slow processor can cause the controller to miss its own deadline. If you process data on behalf of clients, your contract should commit you to a notification window well inside 72 hours (24 hours is a common standard) precisely so your client has time left to meet theirs.

8. Building an Incident Response Plan

An incident response plan turns a breach from a crisis improvised in real time into a rehearsed procedure. At minimum, a workable plan should define:

1. **Who is on the response team** — typically a technical lead, a decision-maker with authority to notify the regulator, and (if you have one) your DPO or external privacy counsel. Name actual people, with backups, not just job titles.
2. **How an incident gets reported internally** — a single point of contact (an email alias or ticketing category) that any employee can use immediately, with no filtering through management first.
3. **A triage step** — is this actually a personal data breach under Article 4(12)? Not every security incident involves personal data, and not every IT problem is a security incident; this step avoids both under- and over-reporting.
4. **A containment step** — how to stop ongoing unauthorised access or data loss (revoke credentials, isolate systems, patch the vulnerability) before full root-cause analysis is complete.
5. **A risk assessment step** — likelihood and severity of harm to individuals, which determines whether Article 33 notification is mandatory and whether Article 34 individual notification is also required.
6. **A notification step** — pre-drafted templates for the supervisory authority notification and, separately, for a plain-language individual notification, so the 72-hour window is not spent starting from a blank page.
7. **A post-incident review** — what allowed the breach, what the fix is, and whether it should feed back into your Article 32 measures or your Article 35 DPIA process.

Test the plan before you need it. A short tabletop exercise once a year — walking through a realistic scenario (a phishing-compromised admin account, a lost company laptop) with the actual response team — reliably surfaces gaps that look fine on paper: nobody knows the regulator's notification portal, the "24/7 emergency contact" is a personal mobile number that changed jobs eighteen months ago, or the plan assumes IT access that the response team doesn't actually have on a weekend.

9. Vendor and Processor Security (Article 28)

Your Article 32 obligations do not stop at your own infrastructure. If a payroll provider, email marketing platform, cloud host, analytics tool, or support-ticket system processes personal data on your behalf, you remain accountable for that data's security as the controller — and Article 28 requires that you only use processors providing "sufficient guarantees" to implement appropriate technical and organisational measures.

Before signing with a new vendor:

- Confirm a **Data Processing Agreement (DPA)** is in place, covering the Article 28(3) mandatory terms — processing only on documented instructions, confidentiality commitments, security measures, sub-processor rules, breach notification obligations to you, and deletion/return of data at contract end.
- Ask what security certifications or audits they hold (SOC 2 Type II and ISO 27001 are the most common credible signals) — and treat the absence of any as a reason to look closer, not necessarily a dealbreaker for a small vendor, but a flag.
- Confirm their sub-processor list and where data is actually hosted, especially for any transfer outside the EU/EEA (this affects your Chapter V international-transfer obligations as well as security).
- Check their breach notification commitment to you — Article 33(2) requires processors to notify controllers "without undue delay," but a contract that specifies a concrete window (24–48 hours) is far more useful in practice than the bare statutory language.

On an ongoing basis: review critical vendors' security posture at renewal, not just at onboarding — a vendor's security can degrade after acquisition, staff turnover, or simple neglect, and your exposure doesn't pause because the contract was signed two years ago. Maintain a vendor inventory (who has access to what personal data, and why) as a living document, not a one-time exercise — this is also the fastest way to answer "who did we need to notify" if one of those vendors has a breach.

10. Cloud Security for GDPR-Regulated Data

Most businesses now run on cloud infrastructure (AWS, Azure, Google Cloud, or SaaS platforms built on them), which changes the security conversation without removing your Article 32 obligations.

Understand the shared responsibility model. Cloud providers secure the underlying infrastructure — physical data centres, network, hypervisor — but you remain responsible for how you configure what you build on top of it: access permissions, encryption settings, network rules, and application-level security. The large majority of publicly reported "cloud breaches" are not failures of the cloud provider's infrastructure; they are misconfigurations on the customer side — a storage bucket left open, an API left unauthenticated, an admin console exposed to the public internet.

Practical baseline for cloud-hosted personal data:

- Encrypt storage and databases at rest (usually a checkbox at provisioning, as noted in Section 3).
- Default all storage (buckets, blobs, file shares) to private, and treat any request to make something public as requiring explicit, documented justification.
- Restrict administrative consoles and databases to known IP ranges or a VPN, rather than leaving them reachable from the open internet.
- Enable your cloud provider's built-in security posture tools (AWS Security Hub, Azure Security Center, Google Security Command Center) — these will flag common misconfigurations automatically, often for free at a basic tier.
- Keep an inventory of what regions your data is actually stored and processed in — cloud consoles make it easy to spin up resources in a new region without noticing, which can silently create an unintended international transfer.

11. AI Systems and Data Security

Generative AI and machine learning tools introduce security considerations that Article 32 already covers in principle, but that many businesses have not yet mapped to specific practice.

Third-party AI tools and your data. If staff paste customer data into a public AI chatbot to draft a response, summarise a document, or analyse a spreadsheet, that data may be transmitted to, and potentially retained or used for model training by, a third party you have no contract with — this is functionally an uncontrolled, unlogged data transfer, and a real Article 32 gap in most organisations today. A written AI usage policy, specifying which tools are approved for use with personal data (ideally ones with an enterprise agreement that disables training on your data) and which are not, is now a baseline control, not an advanced one.

AI systems you build or deploy. If you build features on top of AI models — a support chatbot, a recommendation engine, a document-processing pipeline — the same Article 32 questions apply as to any other system: is data encrypted in transit to the model provider, is access to the underlying training or fine-tuning data controlled, are prompts and outputs logged in a way that itself creates a new store of personal data requiring its own retention and access controls?

Prompt injection and data leakage. AI systems that combine untrusted input (a user's message) with access to sensitive data (a customer database, internal documents) are a new-ish attack surface — a maliciously crafted input can, in some designs, trick the system into revealing data it should not. Treat any AI system with access to personal data as you would any other privileged application: least-privilege access to only the data it needs, and monitoring for unusual query patterns.

This is a fast-moving area, and it intersects closely with Article 22 (automated decision-making) and the EU AI Act's own risk framework — GDPRGard's AI Human Oversight Checklist covers the decision-making side in more depth; this guide's scope is limited to the data-security dimension.

12. Employee Training and the Human Factor

Technical controls fail quietly when the people operating them don't understand why the controls exist. Phishing remains one of the most common breach entry points precisely because it targets people, not infrastructure — and a well-configured firewall does nothing against an employee who enters their password into a convincing fake login page.

Minimum viable training program:

- Onboarding security briefing for every new employee — what personal data they may handle, what the acceptable-use rules are, and who to contact if something looks wrong.
- An annual refresher — even a short one — covering current phishing patterns, password/MFA hygiene, and the incident-reporting process from Section 8.
- Simulated phishing tests, if resources allow, are one of the highest-value ways to measure whether training is actually working, versus just having been delivered.
- Clear, blame-light reporting culture — the earlier a mistaken click or lost device is reported, the more the containment step in your incident response plan can actually contain. An employee who hides a mistake for a week out of fear of consequences turns a minor incident into a major one.

Article 32's "organisational measures" language exists specifically to capture this: encryption and access control are necessary but not sufficient if the people who use the systems were never told how to use them safely.

13. What Regulators Have Actually Fined For

Enforcement data across EU/EEA supervisory authorities shows a consistent pattern in Article 32-related fines: they rarely punish the absence of an exotic control, and rarely turn on genuinely novel attack techniques. They punish the absence of measures that have been considered baseline good practice for years. Recurring findings include:

- **Passwords stored in plaintext or weakly hashed**, discovered after a breach exposed them directly.
- **No multi-factor authentication on administrative or externally-facing accounts**, cited as a failure to implement "state of the art" measures given how widely available and low-cost MFA is.
- **Unpatched, known vulnerabilities** exploited months after a patch was publicly available — the finding is not that the vulnerability existed, but that a fix existed and wasn't applied.
- **Publicly accessible cloud storage** containing personal data, due to a default or misconfigured permission setting.
- **Excessive data retention** — personal data breached long after it should have been deleted under the organisation's own stated retention policy, compounding the severity of the incident.
- **Late or incomplete breach notification** — treated as a separate, additional finding from the breach itself, because Articles 33 and 34 are independent obligations.
- **Disproportionate access rights** — a breach that was small in principle became large in practice because a compromised low-privilege account had access to far more personal data than its role required.

The throughline is that Article 32 enforcement rewards organisations that can show a considered, documented, risk-based approach — even one that didn't prevent every incident — and penalises organisations where the failure reveals that security was never seriously addressed at all. A breach is not, by itself, proof of non-compliance; the absence of appropriate measures is.

14. 20-Point Security Checklist

A practical starting point for reviewing your own Article 32 posture. Not every item will apply to every organisation — treat it as a discussion prompt, not a pass/fail test.

Encryption

1. ☐ All external connections (web, API, admin) run over TLS 1.2+
2. ☐ Databases and file storage holding personal data are encrypted at rest
3. ☐ Company laptops and mobile devices have full-disk encryption enabled
4. ☐ Passwords are hashed with bcrypt/scrypt/Argon2, never encrypted or stored in plaintext

Access Control 5. ☐ Multi-factor authentication is enforced on all admin and privileged accounts 6. ☐ Access follows least privilege, reviewed at least twice a year 7. ☐ Former employees' access is revoked on their last working day 8. ☐ No shared or generic admin logins remain in active use

Detection & Monitoring 9. ☐ Systems handling personal data produce and retain access logs 10. ☐ There is a defined, known internal process for reporting a suspected incident 11. ☐ Alerts exist for unusual data export or access volume, where technically feasible

Breach Response 12. ☐ A written incident response plan exists, naming actual people and their roles 13. ☐ Notification templates for the supervisory authority and for affected individuals exist in draft form 14. ☐ The plan has been tested (tabletop exercise or real incident) within the last 12 months

Vendors 15. ☐ Every processor handling personal data has a signed DPA in place 16. ☐ Critical vendors' security posture (certifications, breach history) is reviewed at renewal

People & Process 17. ☐ New hires receive a security/data-handling briefing during onboarding 18. ☐ An annual refresher training covers phishing and current threats 19. ☐ A written AI usage policy specifies which tools may and may not be used with personal data 20. ☐ Security requirements are considered at the design stage of any new system or feature, not retrofitted after launch

15. About GDPRGard

This guide is provided free by **GDPRGard.eu** — a platform that helps small and mid-sized businesses meet their GDPR obligations without hiring a full-time compliance team. Our self-serve tools generate GDPR-compliant contracts, breach notification letters, privacy policies, and DPIAs in minutes, and our free resources — checklists, calculators, and guides like this one — are built to be genuinely useful on their own, whether or not you ever become a customer.

Explore the rest of our free tools at gdprgard.eu.
