

Data Classification Standard

Owner	[Security owner role]
Approved by	[Policy approver role]
Effective date	[Effective date]
Next review	[Next review date]
Version	1.0

1. Purpose and scope

This standard defines how [Company name] (the "Company") classifies information by sensitivity and how each class is handled. Classification decides who may access information, where it may be stored, how it is shared, and how it is destroyed.

It applies to all information created, received, stored or processed by the Company in any form, including data in the Company's product, in the production cloud environment, in business applications, on laptops, and on paper. It applies to all employees and contractors.

2. Classification levels

Every data set belongs to exactly one of the four levels below. When a data set mixes levels, it takes the highest level it contains. Unlabeled information is treated as Internal until it is classified.

Level	Definition	Examples
Restricted	Information whose disclosure would cause severe harm to customers or the Company, create legal or contractual liability, or enable an attacker to compromise systems.	Customer data stored in the Company's product; personal data of customers' users; authentication secrets, API keys, encryption keys and production credentials; payment card or bank details; employee government ID numbers and health information.
Confidential	Information intended for a defined group, whose disclosure would harm the Company, its customers or partners.	Source code; system architecture and network diagrams for the production cloud environment; security assessments, penetration test reports and vulnerability details; customer contracts and pricing; employee compensation and performance records; unreleased financials.

Level	Definition	Examples
Internal	Information for use by personnel in the ordinary course of business, whose disclosure would cause limited harm.	Internal policies and procedures; team wikis and meeting notes; project plans and roadmaps without customer names; internal announcements; the org chart.
Public	Information approved for release to anyone, whose disclosure causes no harm.	Marketing website content; published documentation for the Company's product; published blog posts and press releases; public job postings; the published privacy notice.

3. Handling rules

Personnel handle information according to the rules for its level. A rule for a lower level applies to every higher level unless a stricter rule replaces it.

Control	Restricted	Confidential	Internal	Public
Access	Named individuals with a documented business need, approved by the data owner and reviewed quarterly.	Role-based access for teams with a business need, reviewed at least every 6 months.	All personnel.	Anyone.
Storage	Only in approved production data stores or an approved secrets manager. Never on laptops, removable media, personal accounts or in source code.	Company-managed systems and Company accounts only. Not on personal devices or accounts.	Company-managed systems and Company accounts.	No restriction.
Sharing	Only with authorized recipients under a signed contract or data processing agreement, through an approved encrypted channel. Never by email attachment or chat.	Inside the Company on a need-to-know basis. Externally only under a signed NDA or contract.	Inside the Company. Externally only with manager approval.	Freely, once approved for release.

Control	Restricted	Confidential	Internal	Public
Encryption in transit	TLS 1.2 or higher on every connection, including internal service traffic.	TLS 1.2 or higher.	TLS 1.2 or higher for network transfer.	TLS on public websites.
Encryption at rest	AES-256 or the provider equivalent, with keys held in a managed key service and access to keys logged.	Encrypted at rest by the storage provider; full-disk encryption on laptops.	Full-disk encryption on laptops.	Not required.
Retention	Kept only as long as the customer contract, law or documented business need requires, then deleted.	Kept for the period set in the retention schedule or contract.	Kept while useful to the business.	Kept while current.
Disposal	Deleted using provider deletion tools or cryptographic erasure; media wiped to NIST SP 800-88 or physically destroyed; deletion recorded.	Secure deletion; devices wiped before reuse or disposal; paper shredded.	Standard deletion; paper shredded.	No restriction.

Restricted data is never used in development, testing or demo environments unless it has been irreversibly anonymized, and never pasted into public AI tools or unapproved third-party services.

4. Roles and responsibilities

Role	Held by	Responsibilities
Standard owner	[Security owner role]	Maintains this standard, approves exceptions, and confirms that controls for Restricted data are in place.
Data owner	[Security owner role], unless a data set has a named owner in the data inventory	Assigns the level to each data set, approves access to Restricted and Confidential data, and takes part in access reviews.
System custodian	[Engineering owner role]	Implements storage, encryption, logging, backup and deletion controls in the production cloud environment and the Company's product to match the level.

Role	Held by	Responsibilities
Endpoint custodian	[IT owner role]	Enforces full-disk encryption, screen lock and approved storage on Company devices.
All personnel	Everyone	Know the level of the information they handle, follow the handling rules, and report suspected mishandling to the [Security owner role] immediately.

5. Labeling

- Documents and files classified Restricted or Confidential show the level in the header, footer, title or file name.
- Shared folders, drives, repositories and data stores holding Restricted or Confidential data show the level in their name or description.
- Databases, buckets and tables holding Restricted data are tagged with the level in the cloud console or infrastructure code.
- Internal information does not need a label. Public information is labeled only when it could be mistaken for a higher level.
- Anyone who receives unlabeled information that appears to be Restricted or Confidential treats it at that level and tells the data owner.

6. Classifying a new data set

Before a new data set is collected, stored or shared, including data brought in by a new feature, integration or vendor, the person introducing it follows these steps:

1. Describe the data set: what it contains, where it comes from, where it will be stored, and who needs access.
2. Check whether it contains any Restricted item listed in section 2. If it does, the whole data set is Restricted.
3. If it does not, apply the definitions in section 2 in order from Confidential to Public and choose the first level that fits.
4. Confirm the level with the data owner. If the two disagree, the higher level applies until the [Security owner role] decides.
5. Record the data set, its level, its owner and its storage location in the data inventory.
6. Apply the handling rules in section 3 before any real data is loaded.
7. Review the level when the data set changes in content or use, and at least once a year.

Location of the data inventory:

7. Exceptions

A deviation from this standard requires a written exception approved by the [Security owner role] before the deviation begins. Each exception records the data set, the rule not met, the business reason, the compensating controls, and an expiry date no more than 12 months away. The [Security owner role] keeps a log of exceptions and reviews it at each annual review.

8. Enforcement and review

Mishandling of Restricted or Confidential information is handled as a security incident under the Incident Response Plan and may lead to disciplinary action under the Code of Conduct. The [Security owner role] reviews this standard at least once a year and after any significant change to the Company's product or the data it processes. The [Policy approver role] approves each new version.

Source

Source: free template from www.polaralabs.com/templates/data-classification-standard (CC BY 4.0).