International Journal of AI, BigData, Computational and Management Studies
Noble Scholar Research Group | Volume 6, Issue 1, PP 227-231, 2025
ISSN: 3050-9416 | https://doi.org/10.63282/3050-9416.IJAIBDCMS-V6I1P123

Original Article

Privacy-By-Design Engineering Under GDPR and CCPA: Practical Patterns for Cross-Border Data Handling In Cloud-Based Applications

Sri Gantikota
Senior Software Engineer, San Diego, California 92101, USA.

Received On: 06/03/2025 ***Revised On: 25/03/2025*** ***Accepted On: 27/03/2025*** ***Published On: 31/03/2025***

Abstract: *Article 25 of the European Union General Data Protection Regulation imposes a legal obligation on data controllers to implement data protection by design and by default. The California Consumer Privacy Act, with its 2020 amendment by the California Privacy Rights Act, reaches a similar outcome through a different mechanism, requiring clear notice, opt-out controls for sale or sharing of personal information, and limits on the processing of sensitive personal information. Together with the wave of comparable regulations enacted in other US states and other jurisdictions since 2020, these requirements have transformed privacy from a compliance footnote into a first-order engineering concern. This paper presents practical patterns for implementing privacy-by-design in cloud-based applications that handle personal data across jurisdictions. The patterns cover data minimization at the schema and API level, purpose limitation through tag-based access control, storage limitation through automated retention enforcement, consent capture and propagation, data subject access request fulfillment, breach notification readiness, and the architectural choices that determine whether cross-border data flows are tractable to govern. Each pattern is presented with the legal rationale that motivates it, the engineering work required to implement it, and the residual risks that remain. The patterns are intended to be tractable for engineering teams without specialized privacy expertise, while preserving the rigor that the regulations require.*

Keywords: *GDPR, CCPA, CPRA, Privacy By Design, Data Protection By Default, Article 25, Data Minimization, Purpose Limitation, Storage Limitation, Consent Management, Data Subject Access Request, Cross-Border Data Transfer, Cloud Applications.*

1. Introduction

The European Union General Data Protection Regulation, in force since May 2018, imposes broad obligations on organizations that process personal data of European Union residents. Article 25 of the regulation specifically requires data protection by design and by default. The data controller must implement appropriate technical and organizational measures both at the time of determining the means of processing and at the time of the processing itself, integrating data protection into the architecture of the system. The default settings of a product must process only the personal data necessary for each specific purpose.

The California Consumer Privacy Act, in force since January 2020 and amended by the California Privacy Rights Act effective January 2023, reaches similar outcomes through different mechanisms. The CCPA and its CPRA amendment require clear notice to consumers about what personal information is collected, opt-out controls for the sale or sharing of personal information, opt-in for sensitive personal information, and rights for consumers to access, delete, and correct their personal information. A growing number of other US states have enacted comparable laws, and other jurisdictions worldwide have enacted regulations that draw on either the GDPR or the CCPA model or both.

This paper presents practical patterns for implementing privacy-by-design in cloud-based applications that operate across these regimes. The patterns are not novel as concepts; the underlying ideas have been part of the privacy literature for decades, going back at least to Ann Cavoukian\'s foundational principles of Privacy by Design articulated in the 1990s. The contribution of the paper is the translation of those concepts into engineering artifacts that a working software team can implement without specialized privacy expertise, while preserving the rigor that the regulations require.

The rest of the paper is organized as follows. Section 2 summarizes the relevant legal obligations. Section 3 covers data minimization patterns. Section 4 covers purpose limitation. Section 5 covers storage limitation and retention. Section 6 covers consent capture and propagation. Section 7 covers data subject access request fulfillment. Section 8 covers breach notification readiness. Section 9 covers cross-border data flow architecture. Section 10 covers limitations and concludes.

2. Legal Background

2.1. GDPR Article 25

Article 25 of the GDPR creates two obligations on the controller. The first, data protection by design, requires the controller to implement appropriate technical and organizational measures designed to implement data-protection principles effectively, considering the state of the art, the cost of implementation, and the nature and risks of the processing. The second, data protection by default, requires the controller to ensure that, by default, only personal data necessary for each specific purpose is processed. The obligation extends to the amount of data collected, the extent of processing, the period of storage, and the accessibility of the data. The European Data Protection Board has published guidance, Guidelines 4/2019, elaborating on how the obligation is to be interpreted in practice.

2.2. CCPA and CPRA

The CCPA, as amended by the CPRA, gives California consumers several rights. The right to know what personal information is collected, used, shared, or sold. The right to delete personal information held about them. The right to opt out of the sale or sharing of their personal information. The right to non-discrimination for exercising these rights. The CPRA additionally creates a category of sensitive personal information with stricter handling requirements and establishes the California Privacy Protection Agency as a dedicated enforcement body. The CCPA does not contain a direct analog of GDPR Article 25, but the practical implementation of the consumer rights it confers requires the same kinds of architectural choices.

2.3. The Comparative Picture

Across the GDPR, the CCPA, and the comparable regulations enacted in other jurisdictions, the common requirements that a cloud-based application must address are similar in shape even where they differ in detail. The application must collect only what is needed, use the collected data only for the purposes for which it was collected, retain it only for as long as necessary, allow the data subject to exercise their rights of access, deletion, and correction, and respond to breaches with appropriate notification. The engineering patterns described in this paper address these common requirements without depending on the specific text of any single regulation.

3. Data Minimization

3.1. Schema-Level Minimization

Schema-level minimization is the practice of designing the data schema to capture only the fields needed for the application\'s documented purposes. The schema is the most durable record of what the application is willing to collect, and a field that is not in the schema cannot be collected accidentally through API drift. The pattern is to start with the smallest schema that supports the application\'s purposes and add fields only when a new purpose is documented. Each field addition is accompanied by an entry in a data inventory that records the purpose, the lawful basis, the retention period, and the data subject categories affected.

3.2. API-Level Minimization

API-level minimization is the practice of designing each API endpoint to return only the fields the caller needs for the documented use case. An endpoint that returns a customer record for display in a customer summary view should not return the customer\'s full payment history, even if the underlying schema includes both. The pattern is to define view-specific response shapes that expose the minimal set of fields and to enforce these shapes at the API layer rather than trusting the caller to ignore extra fields. View-specific responses also limit the impact of compromise of any single endpoint.

3.3. Form-Level Minimization

Form-level minimization is the practice of asking the user only for the data necessary for the immediate transaction. A signup form that collects a date of birth needs a documented purpose for that collection; if the purpose is age verification, the form can collect a boolean attestation of age rather than the full date of birth, depending on the precision required. The pattern requires product designers and engineers to interrogate each form field rather than carrying forward fields from a previous version of the form without examination.

4. Purpose Limitation

4.1. Tag-Based Access Control

Purpose limitation is the requirement that personal data collected for one purpose not be used for another purpose without an additional lawful basis. The engineering pattern that supports this is tag-based access control, in which each data element is tagged with the purposes for which it may be used, and access requests are matched against the tags. A request to use customer email addresses for marketing must be authorized by the marketing tag on the email address field; if the field carries only the transactional tag, the request is denied at the access layer.

4.2. Audit Logging

Audit logging records each access to personal data with the requested purpose and the granted authorization. The log is the evidentiary record that supports demonstrating compliance with the purpose limitation obligation. The pattern is to log at the access layer rather than at the application layer, so that the log captures all access regardless of which application path triggered it. The log is retained for a period that supports the audit and breach investigation timelines required by the relevant regulations.

5. Storage Limitation and Retention

5.1. Per-Purpose Retention Periods

Storage limitation is the requirement that personal data be retained only for as long as necessary for the purposes for which it was collected. Different purposes carry different retention periods. Transaction records may need to be retained for tax or regulatory reasons for several years. Marketing analytics data may be retained only for a much shorter period. The engineering pattern is to assign a retention period to each data element through the same

tagging mechanism that supports purpose limitation, and to enforce the retention period through automated deletion.

5.2. Automated Deletion

Automated deletion runs on a schedule and removes data elements whose retention period has expired. Deletion is final; the deleted data is not recoverable from production stores. Backups are handled separately under a defined backup retention policy that is itself bounded so that backups do not become an indefinite shadow store of data that has been deleted from production. The pattern includes a deletion log that records what was deleted and when, which is what allows the organization to demonstrate that the deletion happened.

5.3. Soft Delete and Hard Delete

Some data elements need to be soft-deleted before being hard-deleted, for example to support a regulator\'s right to inspect records during an investigation. The engineering pattern is a two-stage deletion in which a soft delete marks the record as deleted and makes it invisible to normal queries, and a hard delete removes the record after a second retention period. The two-stage approach reconciles the storage limitation obligation with the legitimate need to preserve records for investigation in specific contexts.

6. Consent Capture and Propagation

6.1. Granular Consent

Where the lawful basis for processing is consent, the consent must be specific, informed, and revocable. The engineering pattern is granular consent capture, in which the user is asked to consent to each distinct purpose rather than to a single bundled consent for everything. Granular consent is recorded in a consent ledger with the timestamp, the version of the consent text shown, and the specific purposes consented to. The user can revoke any individual consent without revoking the others.

6.2. Consent Propagation

Consent must propagate to every component that processes the data. A user who revokes marketing consent must see marketing processing stop across all systems that received their data, not just in the system in which they revoked the consent. The engineering pattern is a consent event stream that downstream systems subscribe to. A revocation generates an event; each subscriber updates its own state accordingly. The pattern requires the architecture to treat consent state as a first-class shared concept rather than as a per-system local concept.

6.3. Sale and Sharing Opt-Out

The CCPA and CPRA require opt-out from sale or sharing of personal information. The engineering pattern is similar to consent propagation but inverted: the default for a California consumer is that sale and sharing are not authorized, and the opt-out signal flows to every downstream system. The Global Privacy Control signal, supported by some browsers, provides a standardized way for a consumer to express the opt-out at the browser level, which the application must honor.

7. Data Subject Access Request Fulfillment

7.1. Subject Access

The right of access requires the controller to provide, on request, a copy of the personal data being processed about the data subject, together with information about the processing. The engineering pattern is a subject access workflow that, on request, gathers the data subject\'s personal data from all the systems that hold it, packages it in a portable format, and delivers it to the data subject. The workflow has to be triggerable through a self-service path that does not require a privacy team intervention for routine requests, with escalation to the privacy team for the cases that require judgment.

7.2. Deletion Requests

The right to deletion requires the controller to delete the personal data on request, subject to certain exceptions. The engineering pattern is a deletion workflow that propagates the deletion across all the systems that hold the data, records the deletion in the deletion log, and produces a confirmation to the data subject. Exceptions, such as data that must be retained for legal compliance, are documented per data category and surfaced to the data subject as the reason the data was not deleted.

7.3. Correction Requests

The right to correction requires the controller to correct inaccurate personal data on request. The engineering pattern is a correction workflow that updates the data in the system of record and propagates the update to dependent systems. The pattern depends on a clear designation of the system of record for each data element, which is itself an architectural property that the data inventory should establish.

8. Breach Notification Readiness

8.1. Detection

Breach notification is required under both GDPR and several US state laws, with timelines that begin running when the breach is detected. The engineering pattern is detection tooling, monitoring, and logging configured so that a breach is likely to be detected promptly, and an incident response process that begins the notification timer at the right moment. The audit logging described in Section 4.2 supports breach detection by providing the access record that the incident response team will examine.

8.2. Scope Determination

When a breach is detected, the controller must determine which data subjects are affected. The engineering pattern is a data inventory that supports rapid scope determination: given a compromised system or compromised credential, the inventory can answer which data categories and which data subjects were exposed. The inventory must be current, which requires it to be maintained as part of the normal engineering process rather than as a periodic exercise.

8.3. Notification Generation

The notification itself must reach the affected data subjects through an appropriate channel and must include the information the regulation requires. The engineering pattern

is a notification system that can generate notifications at scale, with the content templated to satisfy the regulatory requirements and the delivery instrumented to support proof of delivery. The system must be exercised periodically through tabletop exercises so that it actually works under the time pressure of a real breach.

9. Cross-Border Data Flow Architecture

9.1. Why Architecture Matters

Cross-border data flows are subject to specific requirements under the GDPR, which restricts transfers of personal data outside the European Economic Area unless one of several mechanisms applies. The engineering pattern is an architecture in which the physical location of stored personal data is a deliberate decision rather than an accident of cloud provider region selection. The decision is recorded in the data inventory and enforced through deployment configuration.

9.2. Regional Data Residency

Regional data residency keeps personal data within a specified geographic region. The engineering pattern is to deploy regional instances of the application stack and to route requests from a region\'s data subjects to the regional instance. The pattern requires the application\'s identity layer to know which region a data subject belongs to and to route accordingly. Cross-region access for legitimate purposes such as global analytics is handled through aggregation that does not move personal data, or through transfer mechanisms that satisfy the regulatory requirements.

9.3. Standard Contractual Clauses and Adequacy

Where personal data must cross borders, the legal mechanisms include the European Commission\'s Standard Contractual Clauses, adequacy decisions for specific jurisdictions, and binding corporate rules. The engineering team\'s role is to support these mechanisms by ensuring that the transfers actually conform to what the legal instruments cover. This is not a technology choice but an operational discipline: the architecture must make it easy to identify what is being transferred where, so that the legal mechanism is verifiable rather than aspirational.

10. Limitations and Conclusion

10.1. Limitations

The patterns in this paper are general and need to be specialized for the actual application context. A consumer-facing application has different consent challenges than a business-to-business application. A multi-tenant cloud service has different residency challenges than a single-tenant deployment. The patterns are starting points for design discussions, not finished solutions.

10.2. The Role of Privacy Engineering as a Discipline

Privacy engineering as a discipline is maturing. The patterns in this paper draw on the practice of privacy engineering as it has developed since the GDPR took effect, and they are intended to be tractable for engineering teams that do not have a dedicated privacy engineer. As the discipline matures further, tooling will absorb more of the work the patterns currently leave to manual implementation. The architectural choices the patterns recommend are intended to remain sound as that evolution proceeds.

10.3. Conclusion

Privacy-by-design is an engineering discipline as much as it is a legal compliance posture. The patterns in this paper, covering data minimization at multiple layers, purpose limitation through tag-based access control, automated retention enforcement, consent capture and propagation, data subject access request fulfillment, breach notification readiness, and cross-border data flow architecture, address the common requirements that the GDPR, the CCPA, and comparable regulations impose. The patterns are tractable for working engineering teams without specialized privacy expertise. Engineering organizations that apply them are likely to find the compliance burden lower and the residual risk smaller than organizations that treat privacy as a separate concern bolted on to an architecture designed without it in mind.

Acknowledgments

This work draws on engineering practice in healthcare and education software development. The author thanks colleagues across multiple organizations whose practical experience with GDPR and CCPA implementation informed this analysis.

Conflicts of Interest

The author declares that there is no conflict of interest concerning the publishing of this paper.

References

1. European Union. General Data Protection Regulation, Regulation (EU) 2016/679. Official Journal of the European Union, L 119, 4 May 2016. https://scholar.google.com/scholar?hl=en&q=General Data Protection Regulation, Regulation (EU) 2016/679
2. European Data Protection Board. Guidelines 4/2019 on Article 25 Data Protection by Design and by Default, Version 2.0, October 2020. https://scholar.google.com/scholar?hl=en&q=Guidelines 4/2019 on Article 25 Data Protection by Design and by Default, Version 2.0, October 2020
3. State of California. California Consumer Privacy Act of 2018, Cal. Civ. Code 1798.100 et seq. https://scholar.google.com/scholar?hl=en&q=California Consumer Privacy Act of 2018, Cal
4. State of California. California Privacy Rights Act of 2020, Proposition 24. https://scholar.google.com/scholar?hl=en&q=California Privacy Rights Act of 2020, Proposition 24
5. California Privacy Protection Agency. CCPA Regulations, Title 11 California Code of Regulations Division 6. https://scholar.google.com/scholar?hl=en&q=CCPA Regulations, Title 11 California Code of Regulations Division 6
6. Cavoukian, A. Privacy by Design: The 7 Foundational Principles. Information and Privacy Commissioner of

Ontario, 2011. https://scholar.google.com/scholar?hl=en&q=Privacy by Design: The 7 Foundational Principles
7. European Commission. Standard Contractual Clauses for the transfer of personal data to third countries, Commission Implementing Decision (EU) 2021/914. https://scholar.google.com/scholar?hl=en&q=Standard Contractual Clauses for the transfer of personal data to third countries, Commission Implementing Decision (EU) 2021/914
8. European Court of Justice. Data Protection Commissioner v. Facebook Ireland and Maximillian Schrems, Case C-311/18, 16 July 2020 (Schrems II). https://scholar.google.com/scholar?hl=en&q=Data Protection Commissioner v
9. Information Commissioner\'s Office. Guide to the General Data Protection Regulation, United Kingdom. https://scholar.google.com/scholar?hl=en&q=Guide to the General Data Protection Regulation, United Kingdom
10. Commission Nationale de l\'Informatique et des Libertes. Recommendations on cookies and other trackers, France. https://scholar.google.com/scholar?hl=en&q=Recommendations on cookies and other trackers, France
11. National Institute of Standards and Technology. Privacy Framework, Version 1.0, NIST, January 2020. https://scholar.google.com/scholar?hl=en&q=Privacy Framework, Version 1.0, NIST, January 2020
12. International Organization for Standardization. ISO/IEC 27701 Privacy Information Management. https://scholar.google.com/scholar?hl=en&q=ISO/IEC 27701 Privacy Information Management
13. World Wide Web Consortium. Global Privacy Control specification, draft. https://scholar.google.com/scholar?hl=en&q=Global Privacy Control specification, draft
14. Hoepman, J. Privacy Design Strategies. IFIP International Information Security Conference, 2014. https://scholar.google.com/scholar?hl=en&q=Privacy Design Strategies
15. Spiekermann, S. and Cranor, L. Engineering Privacy. IEEE Transactions on Software Engineering, 35(1), 67 to 82, 2009. https://scholar.google.com/scholar?hl=en&q=and Cranor, L
16. Organisation for Economic Co-operation and Development. OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, revised 2013. https://scholar.google.com/scholar?hl=en&q=OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, revised 2013
17. Brazilian National Congress. Lei Geral de Protecao de Dados Pessoais, Law No. 13,709/2018. https://scholar.google.com/scholar?hl=en&q=Lei Geral de Protecao de Dados Pessoais, Law No
18. Government of Canada. Personal Information Protection and Electronic Documents Act, S.C. 2000 c. 5. https://scholar.google.com/scholar?hl=en&q=Personal Information Protection and Electronic Documents Act, S.C
19. Virginia General Assembly. Consumer Data Protection Act, Va. Code 59.1-575 et seq. https://scholar.google.com/scholar?hl=en&q=Consumer Data Protection Act, Va