ConStand

Specification for managing data protection consents
and data subject rights in the digital sphere

Version History

Version number	Date of completion	Adapted by
1.0	15th August 2025	Valentin Rupp (VR), Max von Grafenstein (MvG)
1.1	30th January 2026	VR
1.2	20th February 2026	MvG
1.3	08th April 2026	VR

Table of Contents

1. Introduction

This section is non-normative

This document sets out the full specification for the consent process. It consists of three building blocks—the technical specification, the UX and UI design specification, and the legal specification—and together they form a comprehensive standard for obtaining and managing informed consent for the processing of personal data in digital environments using consent agents. The combined technical, visual, and legal framework is designed to enable genuinely informed consent and, for the first time, establish a level playing field by fostering fair competition that supports data protection–friendly actors. ConStand also serves as a functional interface to further European initiatives, such as the EUDI Wallet, Data Intermediation Services under the Data Governance Act, as well as agentic AI.

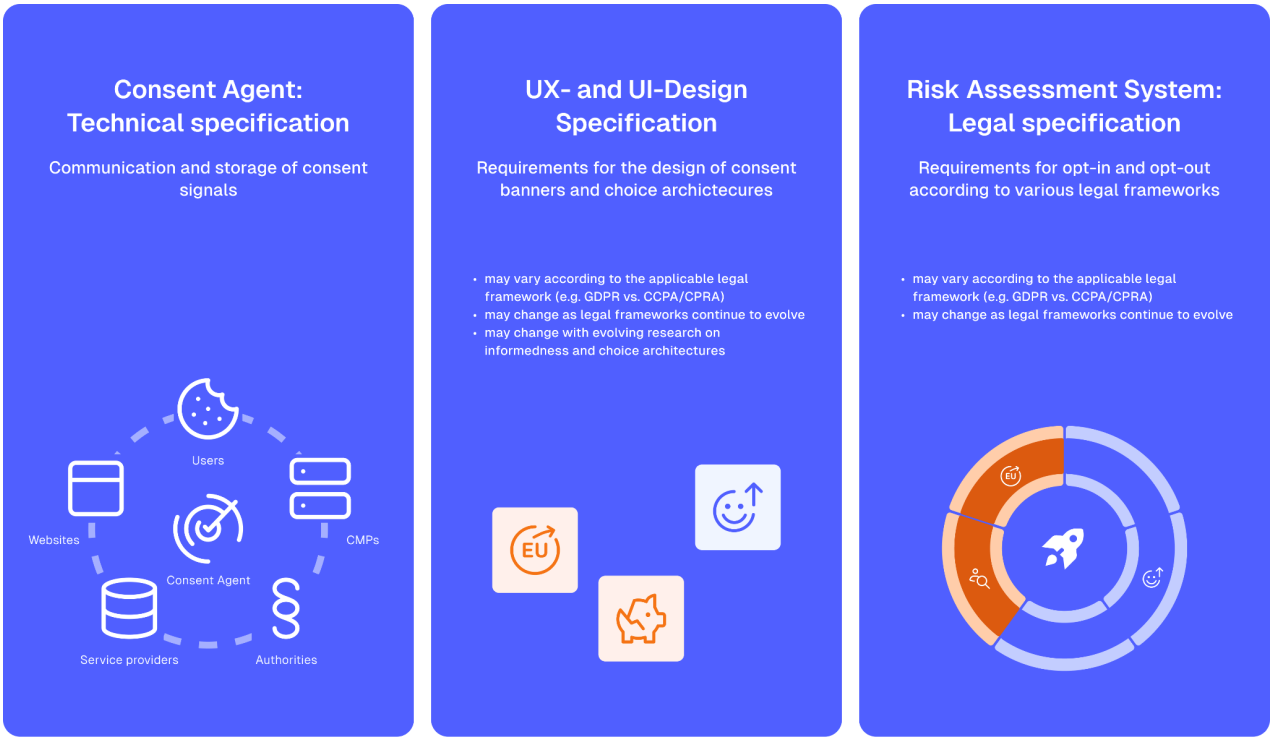

1.1 Functional Overview

This specification defines a mechanism for expressing and managing user decisions regarding the processing of personal data under the European Union's data protection regulations and comparable frameworks outside the EU. The specification establishes standardised rules for the communication of these decisions between multiple stakeholders and components, enabling users to

- provide, refuse or withdraw **consent** (opt-in),
- **object** to specific processing purposes (opt-out),
- or invoke further data subject rights.

In doing so, the specification aims to achieve the following functions:

- Resolving the problem of **consent fatigue**, whereby consumers are unable to make informed decisions simply because of the sheer frequency of requests;

- enabling **informed decisions** that go beyond addressing consent fatigue to prevent the risk that consumers simply opt in or opt out once, yet still fail to make informed decisions;
- enabling a **competitive advantage** for digital service providers who enjoy a high level of trust in their interactions with consumers, for example through a high standard of data protection.

Managing data protection decisions accordingly, requires not only technical harmonisation of the communication itself, but equally legal and visual harmonisation of purposes, risks and benefits, transparency information and user interfaces. This includes:

1. Standardised processing purposes and transparency information: Establishing a shared taxonomy and structure for specifying the purposes of personal data processing as well as the information presented to users before consenting to the processing (data categories, standardised descriptions of the processing operations, storage locations and retention periods).

2. A standardised risk assessment: Enabling controllers and processors to identify and communicate the fundamental rights risks associated with personal data processing, and to mitigate these risks in line with the current state of the art.

3. Standardised signals and technical interfaces: Ensuring that controllers and processors respect data subjects' decisions regarding the processing of their personal data.

4. Standardised visual interfaces: Empowering data subjects to make fully informed choices and to modify these decisions at any time.

Improvement of the State of the Art
In light of ongoing technological progress and the emergence of new data collection practices and associated risks—across IT security, data protection, and artificial intelligence—this specification aims to advance the state of the art also in data protection consent. It seeks to further the advancement of both technical signalling protocols and user experience and interface design. As consent remains the most important legal basis for processing personal data in the digital sphere, current implementations still lack key components of genuinely informed consent (see 1.2.2).

1.2 Challenge

1.2.1 Legal Background

According to Articles 6(1)(a), 7, and 25(1) GDPR and Recital 32, consent must be a freely given, specific, informed, and unambiguous indication of the data subject's wishes, expressed by a clear affirmative act. Data subjects must be able to withdraw their consent as easily as they gave it. According to Article 21(5) and Article 6(1) lit. f GDPR, data subjects may also exercise their right to object by automated means using technical specifications. According to Article 25 GDPR, service providers must implement these intervention mechanisms in a way that enable data subjects to effectively control the risks to their fundamental rights. When implementing these mechanisms, providers must ensure that these mechanisms are **effective** and must verify effectiveness empirically; providers must also consider the so-called **state of the art**, i.e. the most effective implementation available on the market.[1] The dynamic reference to the state of the art ensures that the mechanisms keep pace with technical developments.[2] In addition, service providers must guarantee security and accountability, according to Articles 32 and 5(2) GDPR. These requirements also apply to the implementation of the other data subject rights pursuant to Art. 15 et seq. GDPR.

[1] EDPB, Guidelines 4/2019 on Article 25 Data Protection by Design and by Default, cip. 14 et seq.
[2] Cf. L. A. Bygrave, 'Data protection by Design and by Default: Deciphering the EU's Legislative Requirements' (2017) 1 OLR 105, 107.

1.2.2 Status Quo

In contrast, most current intervention mechanisms, especially consent implementations in the form of cookie banners do not meet these requirements. Empirical studies show that even “best practice” designs fall short to inform data subjects sufficiently.[3] Data subjects do not understand what they are consenting to or refusing to consent to, or what they are objecting to or not objecting to. Even if they did understand, data subjects are unable to make decisions in line with their preferences due to the current design of the intervention mechanisms.

A key reason for this is that current implementations compress all informational elements of consent into a single interaction within the cookie banner—typically presented in a small, cluttered space at the most inconvenient moment, right when users want to access the service rather than pause to interpret complex data processing operations and their associated risks.

Furthermore, the information provided in this “Blitz” process does not inform data subjects about the significance of their respective decisions, meaning that it is meaningless to them. All this leads to the phenomenon widely known as **consent fatigue**: Users routinely click away banners without informed decisions, which cannot be considered free, informed and unambiguous consent.

Why Do Consent Banners Look the Way They Do?

Beyond the economic incentive for controllers to design consent banners that are so irritating that users simply consent to all processing—rather than making a truly informed choice, as intended by Article 7 GDPR—there is another important reason why consent interfaces look the way they do today.

A common assumption in legal scholarship is that all elements of informed consent must be presented and obtained at a single point in time—specifically, the moment the user decides whether to consent, just before accessing a service.[4] This interpretation, however, leads to the problem mentioned above: in theory, users must read and understand complex, detailed information about data processing each time they use a new service, including its potential benefits and risks. In practice, almost no one does this. Most users simply click “accept” or “decline” without much consideration.

It is therefore important to recognize that neither Article 7 nor Recital 32 GDPR requires all conditions for informed consent to be met simultaneously. They may instead unfold across a sequence of user interactions, as long as, by the moment consent is finally given, all relevant requirements are satisfied (see 1.3 Solution).

In addition, current consent mechanisms suffer from several other shortcomings—from insufficient signal integrity in the delivery and storage of data subject decisions up to the fact that data subjects themselves do not receive any proof of whom they have made decisions about and when.

This situation not only leads to considerable legal uncertainty on both sides: data subjects who do not know to whom they have made which decisions; and service providers who are unsure whether their processing of personal data is based on a legally valid legal basis. The current situation also leads to competitive disadvantages for service providers who incur considerable expense to become legally compliant – and to corresponding competitive advantages for those who do not. Numerous studies

[3] See Grassl, P., Gerber, N., & Grafenstein, M. v. (2024). How Effectively Do Consent Notices Inform Users About the Risks to Their Fundamental Rights? European Data Protection Law Review, 10(1), 96-104. DOI: 10.21552/edpl/2024/1/14.
[4] See Stiemerling; Weiß; Wendehorst, Forschungsgutachten zum Einwilligungsmanagement nach § 26 TTDSG: Studie im Auftrag des Bundesministeriums für Wirtschaft und Energie, 2021.

show that the vast majority of users prefer digital services with a high data protection level. However, due to the current lack of transparency and decision-making inefficiencies, they are unable to choose such services. This not only leads to users failing to appreciate the effort that data protection-friendly services put in, but also to a general mistrust among them even towards data protection-friendly services. Services with lower data protection levels therefore not only benefit from having to put in less effort to achieve the same result, but also drag down the overall level of data protection across the entire market.[5]

1.3 Solution

Summary
The mechanism follows a stepwise learning journey across multiple user interactions to support gradual understanding and decision-making. The consent process begins with the **consent agent**, continues with a **handover notice**, and then progresses to the **consent banner**, optionally complemented by **contextual consent** in some scenarios. Each touchpoint has a distinct role in enabling the data subject to understand the processing activities—especially the associated risks and benefits—and to make an informed decision for a specific controller in a specific context.

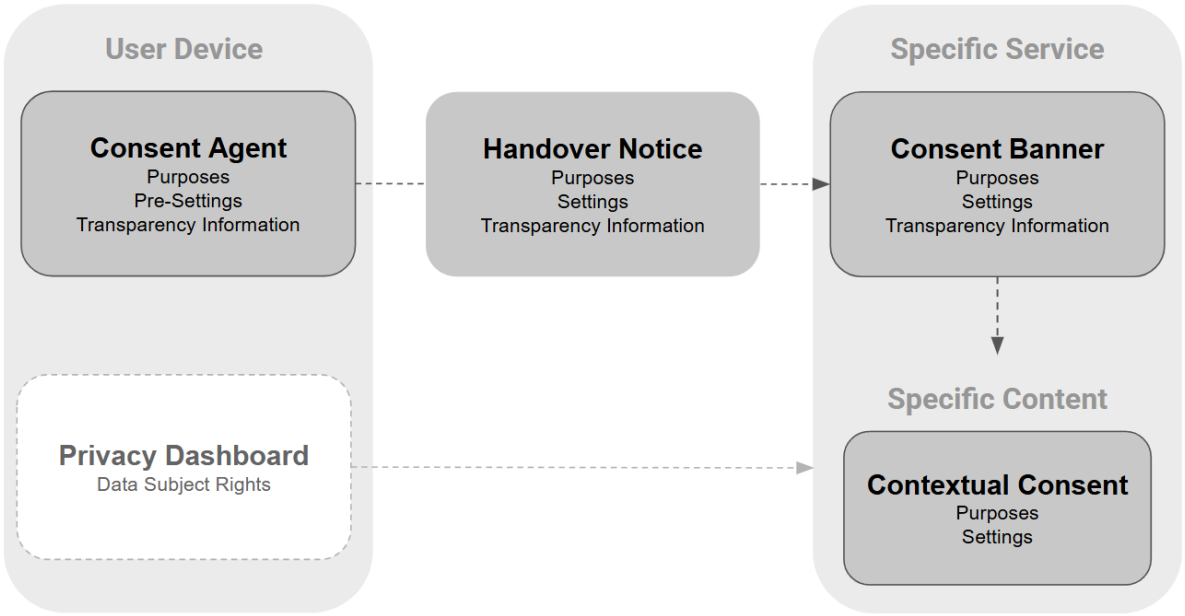

To ensure informed decisions, legal certainty and, on these grounds, competitive advantages for data protection friendly services, the proposed solution conceptualizes consent as a process rather than a single act.

In this model, the conditions for valid consent and similar decisions are distributed across multiple points in time and interaction interfaces. Users can predefine their preferences in a **consent agent**, which transmits these settings to the controller (e.g. websites upon visit). Before data transmission occurs, a **handover notice** (HN) allows users to review or adjust their preferences based on the specific processing circumstances of the site. To avoid consent fatigue, this handover notice disappears automatically after a reasonable period of time and no longer needs to be actively clicked away by users.

[5] This problem is known in economics as the market for lemons, see Akerlof, George A. “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism.” The Quarterly Journal of Economics 84, no. 3 (1970): 488–500. https://doi.org/10.2307/1879431.

The controller employs a **consent banner** to request consent for its required processing purposes and to present all information necessary for the end user to provide informed consent. This consent banner is shown to consent agent users only if they explicitly choose to open the banner to access this information (e.g. through the handover notice or a floating button - see 5.4.3.4). Otherwise the consent decision is being transmitted through the handover notice, without the banner ever appearing.

This system can be supplemented by additional interfaces – especially for obtaining **contextual consent** i.e. directly in the advertising banners or multimedia content in which personal data is being collected.[6] Contextualised consent not only enables data subjects to understand the significance of their decision more directly in the context of their specific usage situation. Contextualised consent also opens up an important design space for service providers to increase user engagement with the extended functionalities of their service. Further interfaces may enable a more usable exercise of data subject rights. They too should be understood in the context of the user's decision-making process and designed accordingly, as they again contribute to a better understanding and control of the data processing and the associated benefits and risks.

Independent Components
Both the consent agent and the consent banner are in principle designed to operate as standalone instances. This specification therefore defines the signals and interfaces that each component, or any provider implementing such a component, must be capable of sending and receiving to properly respect users' decisions. Only this standardisation ensures informed, free, and unambiguous user consent (see 5.).

1.4 Foundational Standards and Specifications

ConStand builds primarily on three categories of standards and specifications that have emerged and evolved over the past nine years, beginning with the Do Not Track (DNT) initiative.

1.4.1 Client Side Signaling Protocols

The first category comprises consent **signaling** mechanisms designed for pre-processing consent decisions in user agents: Do Not Track (launched 2007/2011), Global Privacy Control (GPC, 2020), and Advanced Data Protection Control (ADPC, 2021).

Do Not Track (DNT): An early HTTP header mechanism allowing users to signal opt-out from behavioral tracking. DNT was proposed as a voluntary preference expression without being legally binding. Despite widespread browser implementation, the specification failed due to lack of legal enforcement and voluntary compliance, leading to its discontinuation by W3C.[7]

Global Privacy Control (GPC)[8]: A browser signal enabling users to opt out of data sales/sharing for targeted advertising across sites. Unlike DNT, GPC has legal effects in some jurisdictions with opt-out mandates (CCPA/CPRA, Colorado Privacy Act) and is potentially applicable under GDPR, albeit only for objection proceedings pursuant to Art. 21(5) GDPR (see 5.). It operates as a binary signal at the protocol level, focused on specific legal requests.

---

[6] This is especially relevant for embedding services of third party providers on a website, e.g. via iframes.
[7] See discontinuation notice on https://www.w3.org/2011/tracking-protection/, accessed 17 Feb. 2026; See also Harcourt, Alison, George Christou, and Seamus Simpson, 'The Do Not Track Standard: The Failure of Self-regulation and the Politics of Contestation', Global Standard Setting in Internet Governance (Oxford, 2020; online edn, Oxford Academic, 23 Apr. 2020), https://doi.org/10.1093/oso/9780198841524.003.0007.
[8] Zimmeck et al., https://www.w3.org/TR/gpc/, accessed 1st April 2026.

Advanced Data Protection Control (ADPC)[9]: A comprehensive protocol enabling bidirectional communication between websites and user agents for consent management. ADPC supports purpose-specific consent requests, granular decisions (opt-in and opt-out, i.e. including withdrawal and objections), and aligns explicitly with GDPR/ePrivacy requirements. The protocol allows websites to publish machine-readable consent request lists, to which user agents respond with specific decisions via HTTP headers or JavaScript interfaces.

The standards differ significantly, particularly in the **granularity of user control** they provide:

Standard	Granularity Level
DNT	Binary signal only: track (0), do not track (1), or unset. No purpose or processing-specific distinction possible.
GPC	Binary at protocol level (signal present or absent). Scoped to specific legal request: 'do not sell/share my personal information' and similar opt-outs. No purpose-level granularity.
ADPC	Highly granular and multi-dimensional. Supports purpose-specific consent and refusals, plus multiple action types (consent, refuse, withdraw, object). Controllers can define custom purpose lists or use standardized taxonomies. Allows combined signals (e.g., 'reject all except purposes X and Y').

ConStand builds on ADPC as a purpose-specific framework supporting multiple action types (consent, refuse, withdraw, object) in the following ways:

Additional (meta)data: The signaling is enriched with additional data, such as vendor specific consent, a versioning protocol and timestamps.[10]

Consent receipt: The additional metadata creates detailed consent records across the full consent lifecycle—from creation to withdrawal—leveraging the consent receipts defined in ISO/IEC TS 27560:2023, thus enabling proof of legal compliance (or non-compliance).[11]

Automated risk assessment and versioning: The versioning mechanism integrates with an automated risk assessment, enabling constant monitoring of changing risk profiles per individual consent decision.[12]

UI/UX Requirements: Formulation of clear requirements regarding the exchange and display of transparency information, required for obtaining truly **informed** consent (e.g. through the communication of risks and benefits connected to the intended data processing).[13]

Data minimisation and security requirements: Requirements regarding signal integrity, data storage and data minimisation.[14]

1.4.2 Consent Records

The second category addresses the standardisation of **consent records**, which are created on the basis of information received via consent signaling—including metadata for transparency, accountability, and proof of (non-)compliance. Key here is ISO/IEC TS 27560:2023 (Consent Records

[9] Human et al., https://www.dataprotectioncontrol.org/spec/, accessed 1st April 2026.
[10] For a complete overview of the consent records defined in this specification, see 4.1.2.
[11] See 3.2.1.
[12] See 5.
[13] See 5.
[14] See 6.

and Receipts), which leverages the Data Privacy Vocabulary (DPV) as a standardized RDF-based vocabulary for expressing privacy concepts like consents, purposes, and statuses (see below).

ISO/IEC TS 27560:2023 – Consent Records and Receipts: A technical specification defining an interoperable, machine-readable structure for recording, storing, and exchanging consent information (consent records). The standard specifies four main sections: Record Metadata (identifiers, timestamps), Parties (data subject, controller, third parties), Processing (purposes, data categories, legal basis), and Consent Management (status, validity, withdrawal mechanisms). ISO 27560 supports the full consent lifecycle from creation through withdrawal, enables exchange of consent information between entities via receipts, and aligns with GDPR requirements. It leverages the Data Privacy Vocabulary (DPV) as a standardized RDF-based vocabulary for expressing privacy concepts like consents, purposes, and statuses (see below).

Example for consent receipt (JSON file):


```
",
  "@type": "ConsentReceipt",
  "id": "urn:uuid:550e8400-e29b-41d4-a716-446655440000",
  "issuedAt": "2026-02-13T14:30:00Z",
  "controller": {
    "@type": "DataController",
    "name": "Beispiel GmbH",
    "contact": "datenschutz@beispiel.de"
  },
  "subject": "did:example:user123",
  "consentRecord": {
    "@type": "ConsentRecord",
    "hasPurpose": "dpv:Marketing",
    "hasPersonalData": [
      "dpv:EmailAddress",
      "dpv:Name"
    ],
    "hasProcessing": "dpv:Collection",
    "hasLegalBasis": "dpv:Consent",
    "isWithdrawalPossible": true,
    "withdrawalMethod": "https://beispiel.de"
  },
  "proof": {
    "type": "DataIntegrityProof",
    "cryptosuite": "eddsa-jcs-2022",
    "created": "2026-02-13T14:30:05Z",
    "proofPurpose": "assertionMethod",
    "proofValue": "zQeVbY4oey5q2M3XK..."
  }
}
```


ConStand leverages this specification for creating comprehensive consent records—including data subjects' granular decisions on specific processing purposes, recipients, and transparency details, plus essential metadata—in both complete and data-minimized versions. This ensures full accountability and proof of (non-)compliance for controllers and data subjects, while adhering to data minimization by storing only the information necessary to demonstrate compliance.

1.4.3 Metadata & Vocabulary Specifications

These provide structured, machine-readable languages for describing privacy concepts, data processing activities, and content categorization.

W3C Data Privacy Vocabulary (DPV): A comprehensive RDF/SKOS-based vocabulary providing machine-readable metadata representation for personal data processing. DPV includes extensive taxonomies covering purposes (Marketing, Service Provision), processing operations (Collect, Store, Share), personal data categories, technical and organizational measures, legal bases (including consent types and status), risk assessment, rights exercise, and contextual information. DPV enables representation of privacy notices, consent records, processing activities, risk assessments, and policy specifications.

Note: Compatibility and potential inclusion of the DPV remain to be clarified.

TILT (Transparency Information Language and Toolkit)[15]: A GDPR-aligned formal language and open-source toolkit for representing and processing machine-readable transparency information about personal data practices. TILT enables structured expression of transparency information such as data purposes, legal bases, recipients, retention periods, and data transfers.

YAPPL (Yet Another Privacy Preference Language)[16]: YAPPL (Yet Another Privacy Preference Language) is a machine-readable policy language for expressing granular privacy preferences, such as permitting or excluding specific purposes (e.g., "statistics" but not "commercial"), data recipients (e.g., "wikimedia" only), data transformations (e.g. aggregation), and time-bound conditions (e.g. valid_from and exp_date).

Unlike protocol-driven signals like GPC or ADPC, YAPPL is designed for client-side autonomy: preferences are serialized as structured JSON rules, stored locally, and evaluated via JavaScript against a site's TILT transparency document to automate banner rendering and toggle enforcement (on/off states).

How do TILT and YaPPL work together?
Websites publish their privacy policy as a TILT JSON document, loaded via JS. The toolkit parses TILT to render dynamic banners showing purposes/categories, allowing users to toggle preferences. User choices are then encoded as YAPPL JSON (e.g., consents per purpose) and saved to localStorage for persistence on return visits.

1.4.4 Transparency & Consent Framework (TCF)

Among the specifications introduced above, the TCF occupies a unique niche as the dominant industry-led protocol for ad-tech interoperability. The TCF (developed by the IAB Europe) incorporates signaling, consent record and vocabulary standardization into one large industry framework for managing user consent and transparency for online advertising under the GDPR and ePrivacy Directive. The TCF standardizes how publishers, advertisers, and consent management platforms (CMPs) collect, store, and communicate users' consent choices and legitimate interest objections through a structured consent string. It defines metadata for purposes, features, and vendors, enabling sharing of consent information across the ad-tech ecosystem.

---

[15] Grünewald, Elias & Pallas, Frank. (2021). TILT: A GDPR-Aligned Transparency Information Language and Toolkit for Practical Privacy Engineering. 636-646. 10.1145/3442188.3445925.
[16] Ulbricht, MR., Pallas, F. (2018). YaPPL - A Lightweight Privacy Preference Language for Legally Sufficient and Automated Consent Provision in IoT Scenarios. In: Garcia-Alfaro, J., Herrera-Joancomartí, J., Livraga, G., Rios, R. (eds) Data Privacy Management, Cryptocurrencies and Blockchain Technology. DPM CBT 2018 2018. Lecture Notes in Computer Science(), vol 11025. Springer, Cham. https://doi.org/10.1007/978-3-030-00305-0_23

While the TCF enables large-scale interoperability and thus represents the most comprehensive existing specification, its design primarily reflects industry-driven consent signalling rather than genuinely user-centric control. In this respect, its fundamental approach differs significantly from the aforementioned DNT, GPC, and ADPC initiatives. Most notably, the process for capturing user decisions does not originate in the user agent but is instead initiated separately by each individual website. This perpetuates key flaws in current consent mechanisms, such as intransparency and consent fatigue.

ConStand aims to bridge the gap between operational practicality for data controllers and effective, transparent user control. While designed for integration with existing systems like TCF—particularly in processing purposes and vendor-specific consent structures (see 3.)—it extends these with state-of-the-art features essential for legally valid, informed consent under the GDPR. This includes especially improvements regarding data minimisation, signal integrity and a presentation of information that is much easier for laypeople to understand.

Additional improvements over the current TCF include enhancements for a much more user-friendly exercise of data subject rights (in particular, data access, correction and deletion) and significantly strengthened governance mechanisms that can be linked, for example, via certification programmes in accordance with Art. 42 et seq. GDPR. These additional mechanisms are to be included in future components and are so far not part of this specification.

Overall, these enhancements not only enable fully legally compliant data processing, even for personalised advertising. Rather, the improvements enable service providers to leverage a high level of data protection as a competitive advantage, thereby providing important incentives to develop and deploy technologies that pose ever lower data protection risks, i.e. continuously enhancing the state of the art according to Articles 25 and 32 GDPR.

Standard	Granularity Level
TCF	• Supports purpose-level consent and objections (for up to 10 standardized purposes and several features/vendors). • Vendor-specific: Allows consent for or against individual processing partners. • Protocol role: Operates through encoded consent strings, not via browser-level signals, requiring CMP integration and user interaction per context.

1.4.5 Public Runtime APIs

navigator.consent: A proposed browser API that acts as a shared interface between websites' Consent Management Platforms (CMPs) and consent agents such as browser extensions. It lets sites expose structured information about vendors, purposes, and consent state, and allows consent agents to read and update this state on the user's behalf. For a more detailed outlook on how navigator.consent could be integrated into ConStand, see Annex 2.

1.4.6 Other notable standards and specifications

ISO 31700-1:2023 – Privacy by Design for Consumer Goods and Services: An international standard setting requirements for embedding privacy protections throughout the lifecycle of consumer goods and services, from design through end-of-use. It mandates privacy-by-default settings and aids regulatory compliance, a risk assessment, and consumer rights (access, rectification, deletion) without prescribing specific technologies or methodologies. The standard provides a framework for organizational privacy practices rather than technical specifications.

ConStand builds on the data protection risk assessment framework in ISO 31700-1:2023 and specifies those general requirements for data processing in digital services on the basis of consent (Art. 6 (1)(a) GDPR) and legitimate interests (Art. 6 (1)(f) GDPR).

Risk & Benefits Taxonomy: An important taxonomy for risks and benefits arising from the processing of personal data for data subjects, as well as their corresponding representation, was developed by an interdisciplinary research network. This includes the risk and benefit taxonomy, a privacy icons library, and various layout designs for specific usage contexts.[17]

IAB Tech Lab Content Taxonomy: An industry-standard hierarchical categorization system for web content (version 3.0 includes 26 tier-1 categories and 366 tier-2 categories). Primarily used for contextual targeting and brand safety in digital advertising, it provides consistent content description across publishers and advertisers. The taxonomy includes special category data (SCD) extensions to minimize risks of generating sensitive inferences about race, politics, religion, or other protected characteristics from content signals.

W3C Privacy-Preserving Attribution (Level 1): A W3C editor's draft specifying a browser API for aggregating ad performance metrics (conversions from impressions) without individual tracking. The specification uses trusted aggregation services with differential privacy through noise addition, limiting contributions per browser and site to prevent cross-site user recognition. This addresses the specific use case of advertising attribution while preserving user privacy.

1.5 Terms and Definitions

Term	Definition
browser	A client software application that retrieves web resources via HTTP/HTTPS, renders content using HTML/CSS/JS standards, and provides interactive user interfaces for web applications.
cookie banner, consent banner	A visual element that appears on user interfaces, such as websites to inform visitors about the use of cookies, trackers, and similar technologies for collecting personal data, enabling them to opt-in or opt-out of the processing, as required by Art. 6(1)(a), Art. 6(1)(f), Art. 21 GDPR and the ePrivacy Directive.
consent	Any freely given, specific, informed and unambiguous indication of the data subject’s agreement to the processing of personal data for a given purpose

[17] See at Grafenstein, M. v., Jakobi, T., & Stevens, G. (2021). Effective data protection by design through interdisciplinary research methods: The example of effective purpose specification by applying user-Centred UX-design methods. Computer Law & Security Review, 46. DOI: 10.1016/j.clsr.2022.105722; Jakobi, T., Grafenstein, M. v., Smieskol, P., & Stevens, G. (2022). A Taxonomy of user-perceived privacy risks to foster accountability of data-based services. Journal of Responsible Technology, 10, 1-14. DOI: 10.1016/j.jrt.2022.100029; Grafenstein, M. v., Kiefaber, I., Heumüller, J., Rupp, V., Graßl, P., Kolless, O., & Puzst, Z. (2024). Privacy icons as a component of effective transparency and controls under the GDPR: effective data protection by design based on art. 25 GDPR. Computer Law & Security Review, 52. DOI: 10.1016/j.clsr.2023.105924; Smieskol, P., Jakobi, T., & von Grafenstein, M. (2025). From consent to control by closing the feedback loop: Enabling data subjects to directly compare personalized and non-personalized content through an On/Off toggle. Computer Law & Security Review, 59, 1-22. DOI: 10.1016/j.clsr.2025.106186; Gerber, N., Grassl, P., v. Grafenstein, M. v., (in review at Computer Law & Security Review), From Cookie Banners to Consent Agents: A Comparative Study on Informed Consent and Consent Rates.

	relating to him or her.[18]
contextual consent	Consent that is requested in direct connection with a user's interaction with specific content within the service being used. Examples include unlocking a video, a map or a social media link, but also the personalisation of advertising in the context of a specific advertisement displayed, as well as the personalisation of the service or a function of the service when that specific function is being used.
consent agent	A dedicated user agent (e.g. browser extension) that enables individuals to declare general consent preferences for data processing purposes across multiple websites or services. These preferences, including granular choices on purposes, associated risks, benefits, and controllers, are persistently stored locally or in a user-controlled repository.
consent intent	If consent is withdrawn directly in the consent agent, there is currently no way to immediately transmit this decision to the relevant controller, because the consent agent has no direct connection to the controller's consent store API. Consent records are always created by the consent banner (CMP), not by the consent agent. For this reason, agents record a consent intent when users withdraw consent in the agent's history and inform users that this change will be applied once they revisit the controller's service (website). When users return to the website, this consent intent is transmitted to the banner, which then creates an updated consent record on the basis of this `consentIntent`.
consent object	A part of the consent record, which defines for each data data recipient the ● time of consent, ● the ID of the consent record, in which consent was given and/or withdrawn in ● and the consent initiator - i.e. the UI in which consent was given (e.g consent banner "CB" or handover notice "HN").
consent record	A signal that contains one or more individual consent decisions of a user for a specific controller. This signal may exist in different variants, be stored in various locations, and be exchanged between multiple stakeholders.
controller	The entity that provides a service and determines the purposes and means of the processing of personal data or other information stored in the terminal equipment of the user.
handover notice	A handover notice is a visual UI element on websites or apps that transmits (hands-over) pre-configurations for personal data processing from a consent agent to a controller. It allows users to review these pre-configurations and customize them specifically for that controller.
mobile app	A computer program designed to run on a mobile device, such as a smartphone or tablet.
pre-settings	Configurations established by a data subject within a consent agent, specifying consent or refusal for the processing of personal data tied to

[18] Rec. 32 (1) GDPR.

	particular purposes.
processor, data processor	A ‘processor’ means a natural or legal person, public authority, agency or other body which processes data on behalf of the controller.[19]
service	An information society service provided by the controller or a third party service provider (see “third party service”) within which personal data from users is collected (e.g. a website).
third party service	Any tool, software, or other service provided by an entity other than the controller that processes personal data of the user—such as analytics software, personalization models, or cloud storage solutions.
transparency information	All information that needs to be made accessible to the data subject according to Art. 12 ff. GDPR.
user	The person visiting or interacting with the website. This specification uses the word “user” as a term that includes both “data subjects” as defined under Article 4(1) GDPR and “users” as defined in Article 2(a) ePrivacy Directive.
website	The information society service through which the user interacts with a controller. The controller may communicate with the user via the website. A website is delineated by its URL, where any URLs whose origins are schemelessly same site are understood as belonging to the same website. As this specification is intended to standardise consent communication also in other contexts—e.g. mobile apps or other digital interfaces—the term “website” is used in this specification as a general reference encompassing these contexts.

[19] Data Privacy Vocabulary (DPV), Version 2.3“ (W3C DPVCG, 2026). https://w3id.org/dpv#DataProcessor.

2. Components

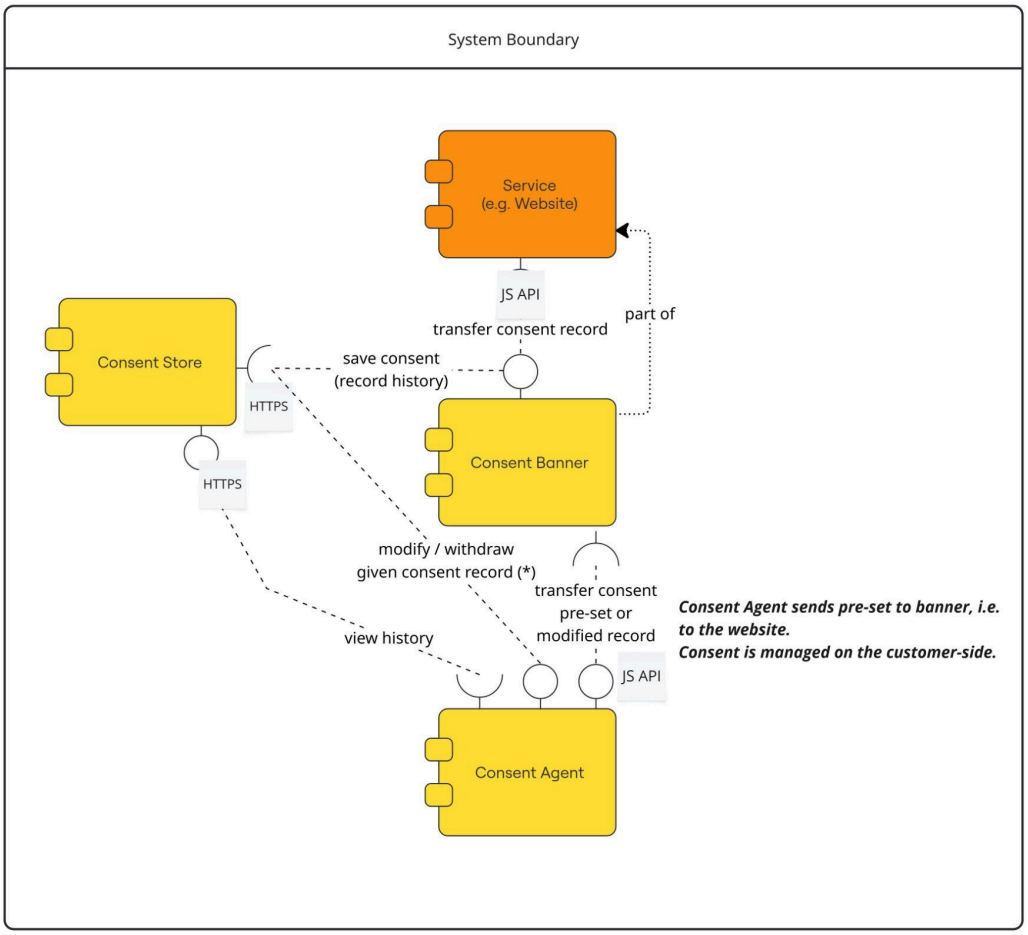

The specification describes a consent process spanning multiple components and stakeholders to enable informed consent for data subjects, facilitate the exercise of data subject rights, and provide all parties with legal proof of compliance (or non-compliance).

The technical components include:

- the consent agent 2.1 (including the handover notice 2.1.1 and the consent history 2.1.2)
- the consent banner 2.2,
- the consent store (including dedicated API) 2.3,
- a windowAPI that enables communication between these components 2.4,
- and contextual consent 2.5.

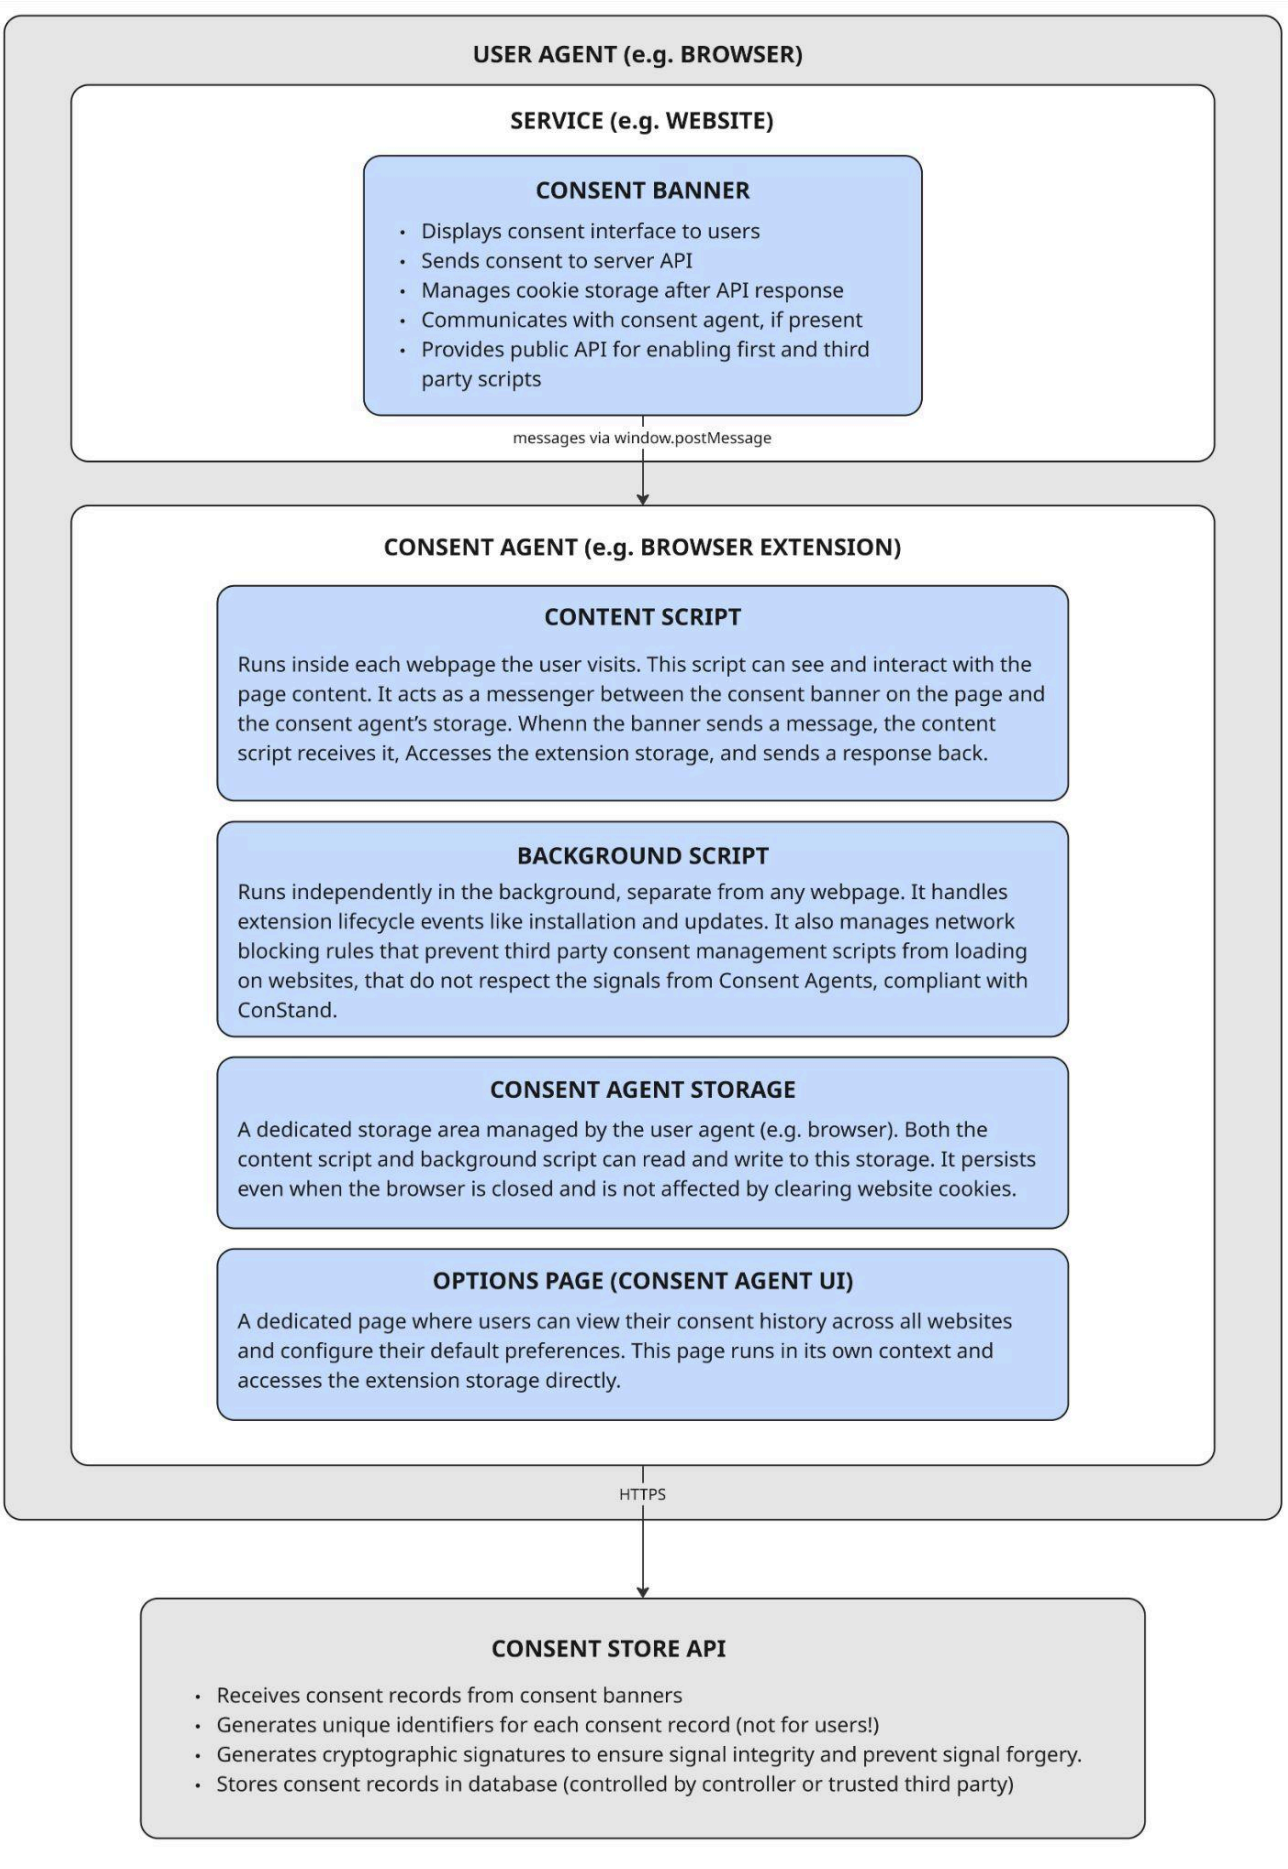
USER AGENT (e.g. BROWSER)
SERVICE (e.g. WEBSITE)
CONSENT BANNER
• Displays consent interface to users
• Sends consent to server API
• Manages cookie storage after API response
• Communicates with consent agent, if present
• Provides public API for enabling first and third party scripts
messages via window.postMessage
CONSENT AGENT (e.g. BROWSER EXTENSION)
CONTENT SCRIPT
Runs inside each webpage the user visits. This script can see and interact with the page content. It acts as a messenger between the consent banner on the page and the consent agent's storage. Whenn the banner sends a message, the content script receives it, Accesses the extension storage, and sends a response back.
BACKGROUND SCRIPT
Runs independently in the background, separate from any webpage. It handles extension lifecycle events like installation and updates. It also manages network blocking rules that prevent third party consent management scripts from loading on websites, that do not respect the signals from Consent Agents, compliant with ConStand.
CONSENT AGENT STORAGE
A dedicated storage area managed by the user agent (e.g. browser). Both the content script and background script can read and write to this storage. It persists even when the browser is closed and is not affected by clearing website cookies.
OPTIONS PAGE (CONSENT AGENT UI)
A dedicated page where users can view their consent history across all websites and configure their default preferences. This page runs in its own context and accesses the extension storage directly.
HTTPS
CONSENT STORE API
• Receives consent records from consent banners
• Generates unique identifiers for each consent record (not for users!)
• Generates cryptographic signatures to ensure signal integrity and prevent signal forgery.
• Stores consent records in database (controlled by controller or trusted third party)

A future version may also include specifications for facilitating and automating the exercise of data subject rights, for example through a **"privacy dashboard"** that allows individuals to manage their rights directly with controllers. This mechanism could build upon the same underlying components and communication channels to support these interactions.

The current specification as well as future components are designed to work independently of any specific operating system (OS) and enable GDPR-aligned consent not only on websites but also in non-web settings such as medical offices, mobile apps, and IoT systems like building information management.

2.1 Consent Agent (incl. Handover Notice and Consent History)

The central component comprises a technical-organisational, visual, and legal infrastructure for a **consent agent**. Through this agent, end users can predefine affirmative or rejecting consent preferences (opt-in) or object to the processing of personal data (opt-out) for specific processing purposes. When making these choices, users are presented with standardized transparency information within the consent agent. This information comprises all elements legally required for informed consent, already linked to specific purposes[20]—though still provided on a general level, meaning not yet tailored to individual controllers but referring to typified processing operations, including typical data categories, storage periods, as well as associated risks and benefits. The information is displayed in a clear, understandable, and transparent manner (e.g., through concise explanations and standardized, tested icons).

This allows data subjects to inform themselves about the intended processing purposes, associated operations, and typical risks or benefits at a point in time when they have the opportunity to reflect on the implications for their fundamental rights—thus enabling them to make a truly informed decision, ideally only once. Via the handover notice, data subjects may adjust their general preferences with respect to the individual situation (i.e. data protection level) of the website visited, though only if they want. Since the handover notice disappears after an appropriate period of time on its own, such an agent-based process does not require the data subjects to actively click cookie banners away on every single website but leaves them in a „lean back" situation where they can actively adjust their preferences but don't have to.

2.1.1 Handover notice

The agent communicates the end user's preferences to controllers (e.g., specific websites visited) via a technical interface as explicit consents or objections. Before any personal data collection or processing occurs, the end user receives a notification—via a dedicated visual **handover notice**—containing a summary about the impending transmission of these consents for the controller's requested processing purposes. In this handover notice, the end user can modify the pre-settings for the specific controller in both directions, either by deselecting purposes that were generally accepted in the consent agent or by selecting purposes that were generally not accepted in the consent agent.

Example 1
A user U sets preferences in a consent agent for predefined processing purposes, agreeing to purposes 1 and 2 while refusing purposes 3 and 4.

[20] A functional system spanning multiple platforms and service providers requires the standardization of purposes, risk assessments, and related terminology (see 3.).

Website operator (controller) C requests consent from visitors for purposes 1–3. When U visits C's website, the consent agent prepares to transmit U's relevant preferences to complete the consenting process. Before transmission occurs, a handover notice appears, listing all purposes requested by C (1–3), clearly indicating U's pre-settings from the agent (affirmative for 1 and 2; refusing for 3), and providing options to modify these on a per-purpose basis for this specific site.

If U takes no action, the affirmative presets for purposes 1 and 2 are finalized as explicit consents, transmitted to C, and stored in the consent store; purpose 3 remains refused as preset. If U instead takes action and, for example, deselects purpose 2 in the notice, only consent for purpose 1 is transmitted and stored, overriding the general preset.

2.1.2 Consent History

The data subject's consent decisions for each visited website are ultimately stored in the consent history, which forms part of the consent agent. This consent history enables users to review and modify their past given and refused consents and can serve as evidence in data protection complaints if a controller has allegedly processed personal data without obtaining informed consent.

2.2 Consent Banner

The second component of the system consists of a user interface for data protection consents (**consent banner**), which can be integrated into the services of controllers (e.g. into the website of a website operator). When using a service that has integrated this consent form, end users of the consent agent only see the handover notice, which they can use to adjust their default settings (see already 1.4.1). For further information on the specific processing operation of a controller, users can open the consent form (e.g. via the handover notice or a floating button placed on the UI of the respective service - see 5.4.3.4).

The consent banner provides all information formally required for informed consent, such as processing purposes, data categories, recipients, storage location, and duration—but now specified for the specific controller. The banner must be both technically compatible with the consent agent via a standardized interface and legally and visually compliant with regard to the information presented to data subjects. In particular, the banner employs identical purpose formulations and communicates purpose-specific risks and benefits consistently with the consent agent.

Attention: If this information is not synchronised in the consent agent on the one hand and the consent form on the other, no transfer or legally effective consent can be assumed because the essential elements of the agreement do not coincide.

Data subjects can view and withdraw their consent at any time, both in the consent history of the consent agent and in the consent form of the respective service. Users without a consent agent are shown the consent form in the conventional manner when they use the service of the controller.

2.3 The Consent Store API (ensuring signal integrity)

The Consent Store API is a server side service that provides permanent storage for consent records (see also 4.2.1). When a user makes a consent decision, the record is sent to this API before being stored locally.

The API performs several important functions:

- It assigns a globally unique identifier to each consent record using UUID version 4 format, ensuring every record can be individually referenced.
- It creates a cryptographic signature using an elliptic curve algorithm that can later verify the record has not been tampered with.
- It stores the record in a database located in an EU data center where it can be retrieved for compliance purposes.

The API returns both a full consent record and a data minimized version, each with their own signature. The full record contains complete consent history, while the minimized version removes withdrawal details for privacy (see further 6. Data Protection and Security Considerations) .

2.4 Public Runtime API (`window.consenter`)

> The Window API defined in this specification is exemplary for the implementation by the consenter agent. It is intended to be replaced in the future by a shared transport and coordination layer, such as navigator.consent (see Annex 2). Achieving this will require additional standardisation work and broad integration by browsers.

In addition to internal message passing, the Consent Banner exposes a public runtime API on the webpage through `window.consenter`. This API allows first party scripts (for example analytics wrappers or tag managers) to read the current consent state and react to consent changes.

The API is initialized by the banner script as soon as it loads:

- `window.consenter.get()`
- `window.consenter.subscribe(callback, identifier)`
- `window.consenter.unsubscribe(identifier)`

The script also dispatches a DOM event `consenter:ready` when this API is available.

2.4.1 API Methods

```
get() Returns the current in-memory consent state in the form:
{
  "status": "pending | no_consent | has_consent",
  "data": { "...ConsentRecord fields..." }
}
```

- `pending` means consent has not yet been resolved in the current page lifecycle.
- `no_consent` means no consent record is available.
- `has_consent` means a consent record is available and has been published.

`subscribe(callback, identifier)` Registers a callback that is called whenever a consent record is published (for example after loading an existing record or after the user saves new choices).

- `callback` receives the full consent record object.
- `identifier` is a unique string used to manage that subscription.
- Returns `"success"` on successful registration.

`unsubscribe(identifier)` Removes a previously registered subscription by identifier. Returns `"success"` when processed.

2.4.2 How It Fits the Existing Flow

`window.consenter` does not replace cookie, extension, or server storage. Instead, it is a runtime integration layer on top of those mechanisms:

1. Consent is retrieved from cookie and/or extension as described above.
2. After consent is resolved, the banner publishes the active record internally.
3. `window.consenter` subscribers are notified so page scripts can immediately adapt behavior (for example enabling or disabling specific services).

When consent is updated by the user, the same publish-notify mechanism runs again after the save flow completes.

2.4.3 Typical Integration Pattern for Website Scripts

1. Wait for `consenter:ready` (or check if `window.consenter` already exists).
2. Call `window.consenter.get()` for immediate state.
3. Register `window.consenter.subscribe(...)` to handle later changes.
4. On teardown/navigation for SPA integrations, call `window.consenter.unsubscribe(...)`.

2.5 Contextual consent

In certain scenarios, controllers may seek to obtain consent directly within the specific usage context of data collection—known as contextual consent. To effectively combat consent fatigue, the data subject's consent agent pre-configurations must be respected in these cases as well.

Important for controllers: Unlike consent obtained via consent banners, contextual consent still allows controllers to request consent more than once. Users may toggle consent on- and off, depending on the specific content they are being displayed in the specific context (such as embedded videos, advertising banners or other content that could be personalised). Consent fatigue is not a concern here, as this consent fatigue is primarily driven by repeatedly appearing banners, which interrupt the flow of users on the website.

In contrast, contextual consent does not contribute to consent fatigue, as it is displayed in direct connection with the use of a specific function or piece of content within the service. The benefits as well as the risks of giving consent are thus intuitively apparent to the user. Of course, this presupposes that contextual consent is implemented in accordance with UX design principles and is not misused to the detriment of the user.

Example 2
User U sets preferences in the consent agent, approving purposes 1–3 while refusing purpose 4. Purpose 3 ("unlock additional website features") specifically authorizes data transmission—such as IP addresses—to third-party providers for loading embedded content like videos, maps, or social media widgets on controller C's site.

Website operator C embeds a maps service from third party provider T via an iframe, which

requires U's IP address to initialize and render properly. When U visits C's site, the consent agent automatically transmits the relevant presets (affirmative for purpose 3) to C, forming valid, informed consent. This consent allows C to directly load the iframe content upon U visiting the website, causing processing of U's IP-Address without further intervention.

The service-specific consent for data processing via T is stored in both the consent store and the consent agent.

Example 3
In this case, C uses the same iframe. However, third-party provider T also collects visitor data to create personalized advertising profiles (purpose 4). Since U's agent presettings explicitly reject purpose 4, the handover notice contains no affirmative consent for this purpose, and none is transmitted to C.

Consequently, C blocks the iframe entirely to prevent unauthorized data collection by T. A visual placeholder or blocked state appears, ensuring GDPR compliance. If U wishes to access the maps despite this, they must override the preset by directly consenting to purpose 4 within the specific iframe context—e.g., via an on-demand consent prompt—unblocking the content and authorizing the necessary data flow.

Increased transparency through contextual consent
Contextual consent has the advantage of making specific processing operations—and their effects in a given situation—transparent to users, thereby enabling them to decide for or against data processing in light of these concrete impacts.

One example is the use of contextual consent directly within advertising iframes. When advertising is shown to a user based on re-targeting (for instance, because they previously added a specific item to the shopping basket in online shop S), users can immediately see how their consent affects their browsing experience and which ads are displayed. If user U disables re-targeting directly in the context of the advertising iframe, the ad would switch to (potentially less relevant) contextual advertising. Conversely, if the user enables re-targeting, they can understand how their behaviour (e.g. adding an item to the shopping cart) influences their experience and shapes their subsequent shopping decisions. This improves users' understanding of both risks and benefits and can reduce the perceived "creepiness" of re-targeting. As re-targeting in practice often entails lower privacy intrusion than profile-based advertising, advertisers can use contextual consent to communicate this more transparently to users. By making the underlying mechanisms visible in the specific context in which they operate, users better understand how the system works, which can alleviate the feeling of being constantly observed.

In addition, contextual consent can be used by advertisers to obtain direct feedback from users. For example, if a user no longer wishes to see re-targeted ads (because they might have already bought the respective item elsewhere or are no longer interested), advertisers can provide a feedback option such as "do not show this ad again," thereby improving the user experience by avoiding annoying, repetitive, or irrelevant ads.

3. Standardised Terminology (esp. purposes, risks and benefits)

A coherent and interoperable consent process across all components relies not only on technical standardisation but also requires standardisation of processing purposes and other transparency

information as well as a consistent framework for determining the risks and benefits associated with each purpose.

Only where

- the purposes used by individual controllers match those presented to consent agent users AND
- the risks and benefits attributed to each purpose in the agent accurately reflect those specified by an individual controller

can data subjects make informed decisions in their consent agent pre-settings already well in advance.

Harmonisation is essential not only for achieving consistent technical communication but also for enhancing understanding among data subjects. The specification and description of processing activities—through the explicit detailing of elements such as data categories, processing locations, purposes, and retention periods—constitute a cornerstone of data protection transparency. In doing so, they provide the principal point of reference for data subjects to evaluate and comprehend the risks involved, enabling them to act accordingly—for instance, by choosing to use or avoid certain services, by granting or withholding consent, and by making well-informed decisions overall.

Example 4
Consent agent user U sets preferences in the consent agent, approving only purpose 1 (“improve the service”). U consents because this purpose restricts controllers to processing personal data solely for aggregated statistics on general user interactions with the website, limiting insights into U's private life.

Controller C receives this consent signal and may process U's personal data based on it only if C has explicitly requested consent for this exact purpose and if it matches U's approval in the agent. C cannot repurpose the data to build individual profiles—e.g., by claiming they are used to “improve the service”—since the purposes (and their associated fundamental rights risks) differ fundamentally, rendering consent invalid for profiling.[21]

3.1 Mapping of processing purposes

While several taxonomies exist for harmonizing processing purposes and transparency information (see 1.4), these taxonomies only partially align with one another. Although full harmonization into a single industry standard would be desirable, achieving this goal is beyond the scope of the present specification.

Instead, this specification defines a base interoperability layer that requires all implementing parties to ensure accurate mapping between different taxonomies when consent information is exchanged between stakeholders.

Given the widespread adoption of the Transparency and Consent Framework (TCF), we propose that its purposes could be used as (one) fundamental reference (see Annex 1). However, controllers, as well as consent agents MAY continue to define their own purpose formulations to better reflect their specific data processing activities. In these cases however, custom purposes would need to remain mappable to the corresponding TCF purposes in order to ensure interoperability across systems. We suggest this mapping be published to allow all stakeholders to review, validate, and, if necessary, correct or improve the alignment.

[21] For this and other cases of non-conformity with this specification, see 10.

Note: Other purpose taxonomies, such as the DPV could equally be used as fundamental reference. The importance lies in the mapping of purposes in order to ensure compatibility across all components and stakeholders. An exemplary mapping for the purposes used in the specification (below) is shown in Annex 1.

Purposes on the consent agent level
At the consent agent level, we propose the following purposes:

- Improve the service
- Unlock additional website features
- Personalize the website
- Support marketing analytics
- Receive marketing offers
- Receive personalized marketing offers
- Customize online ads (non-TCF)

3.2 Mapping of other transparency information

In the same way as for the specification of purposes, additional vocabularies MUST also be used consistently across all components or mapped in accordance with the approach described in 3.1.

3.3 Outlook

Although this specification does not (yet) include a final framework for these purposes, risks and benefits, we propose two things:

1. An open source repository in which processing purposes compatible with this specification are suggested, and agreed upon by all relevant stakeholders and interest groups, esp. service providers, researchers, authorities and data subjects (see standards listed above under 1.4) and

2. A standardised risk-assessment, enabling controllers to specify the risks and benefits—resulting from the processing of personal data for a given purpose—according to their specific processing operation, and inclusion of other data recipients and third party services (see esp. **ISO 31700-1:2023** listed above under 1.4 as well as the data protection risk assessment methodologies developed by the data protection authorities).

It may be reasonable to reference the Data Privacy Vocabulary (DPV) for such a purpose repository. However, we find the vocabulary not very suitable because of the large number of purposes and the somewhat vague categorisation of some purpose types. The DPV's purpose taxonomy also does not consider the risks associated with processing purposes when specifying and categorising them, which we regard as essential. That said, this point is open for discussion.

4. Technical overview

For the entire consent process to function properly, signals and associated metadata must be transmitted between different stakeholders and components and stored in different locations. The information contained in these signals and metadata must follow the principle of data minimisation and thus be limited to what is necessary for achieving the main objectives of the consent process. These objectives include:

- Enabling **data subjects** to make their decision once, without having to repeat it each time they return to the service.
- Enabling data subjects to adjust their preferences
- Allowing **data subjects** to demonstrate to the controller that consent was not given at a specific time or that they had objected to processing.
- Ensuring **data subjects** can change their decisions at any time.
- Enabling **controllers** to demonstrate to data subjects, authorities, and courts that valid and informed consent existed at a given time.
- Enabling **controllers** to change aspects of its data processing (e.g. change a TPP or its configuration) and to inform the data subject accordingly when the data subject is revisiting the website after having already made a choice before.

4.1 Consent record

4.1.1 Structure Overview

A Consent Record captures everything about a user's consent decision at a specific moment in time. It includes the actual choices made for each **purpose** and **third party script**, **metadata** about when and where the consent was given, and **identification information** that links it to a specific configuration of the controller‘s service (e.g. website).

The record structure is designed to support auditing and compliance requirements. The requirements consist of the following components:

1. Preserving the complete history of consent changes,
2. Tracking the source of each decision, and
3. Including version information that allows the controller and user to understand exactly what configuration the user was presented with.

4.1.2 Complete Record Structure

```
ConsentRecord {
    banner_version: "1.0.0"
    domain: "example.com"
    domain_id: "abc123def456ghi78901"
    date: "2024-01-15T10:30:00.000Z"
    consent_record_id: "550e8400-e29b-41d4-a716-446655440000"
    signature: "s1:::MGUCMBglSuTr..."

    trigger: {
        "service-improvement": "1.0.0"
        "additional-features": "0.0.0"
        "personalise-website": "0.0.0"
        ...
    }

    consent: {
        "service-improvement": {
            "matomo": {
                given_at: "2024-01-15T10:30:00.000Z"
                given_in: "550e8400-e29b-41d4-a716-446655440000"
```

```
                consent_initiator: "CB"
            }
            "google-analytics": {
                given_at: "2024-01-15T10:30:00.000Z"
                given_in: "550e8400-e29b-41d4-a716-446655440000"
                withdrawn_at: "2024-02-20T14:45:00.000Z"
                withdrawn_in: "660e8400-e29b-41d4-a716-446655440001"
                consent_initiator: "CB"
            }
        }
        "additional-features": {
            "youtube": {}
        }
        "customise-ads": {}
    }
}
```

The full consent record contains complete information including the entire history of consent changes with all timestamps. This version is needed for comprehensive auditing and demonstrating the complete consent lifecycle.

4.1.3 Data Minimized Record

Not all information contained in a consent record needs to be made available to every party involved in the communication. Therefore, and to comply with the data minimization principle, the system also maintains a data-minimized version of the consent record.

The minimized consent record removes withdrawal information while preserving proof of active consent. When a user has withdrawn consent for a service, the minimized record contains an empty object for that service rather than the detailed withdrawal history. This minimized version is what gets stored in the browser cookie.

```
┌──────────────────────────────────────────────────────────────────────┐
│                    DATA MINIMIZATION EXAMPLE                         │
│                                                                      │
│                                                                      │
│   FULL RECORD (stored on server and in extension):                   │
│   {                                                                  │
│       "google-analytics": {                                          │
│           given_at: "2024-01-15T10:30:00.000Z"                       │
│           given_in: "record-id-1"                                    │
│           withdrawn_at: "2024-02-20T14:45:00.000Z"                   │
│           withdrawn_in: "record-id-2"                                │
│       }                                                              │
│   }                                                                  │
│                                                                      │
│                                                                      │
│   MINIMIZED RECORD (stored in cookie):                               │
│   {                                                                  │
```

```
        "google-analytics": {}
    }


    The minimized record proves that no current consent exists,
    without revealing the withdrawal history
```

Both versions receive independent cryptographic signatures, allowing either version to be verified independently (See 6.4.1).

4.1.4 Field Explanations

banner_version indicates which version of the consent banner configuration was active when the user made their decision. This is important because the available purposes, services, and how they are presented can change over time. Recording the version, documents exactly what the user saw and agreed to (such as storage locations, retention periods and data categories).

domain refers to the service where consent was given, such as the website adress: "example.com" or "shop.example.org".

domain_id is a unique identifier for the service configuration regarding a specific consent banner. While the domain only tells the website address, the domain_id connects to the specific configuration including which purposes, services are available in the specific configuration of the consent banner.

date records when the consent decision was made, formatted as an ISO 8601 timestamp with full precision.

consent_record_id is a globally unique identifier assigned by the server when the record is stored. This identifier uses the UUID version 4 format and ensures every consent record can be individually referenced.

signature is a cryptographic signature created by the server that can be used to verify the record's authenticity. The signature format includes a key identifier followed by the actual signature data.

trigger contains version numbers for each purpose (see 5. Version Management). These numbers are used to determine when users need to be notified about changes or asked to reconsent. The format is three numbers separated by periods, representing different types of changes that might require user attention.

consent contains the actual consent decisions organized by purpose and then by third party service (TPP). For each TPP, we record when consent was given, which record it was given in (consent_record_id), and potentially when it was withdrawn. Empty objects indicate no consent has been given for that service:

```
"youtube": {}
```

Consent intent indicates that a user has withdrawn consent directly in the consent agent and communicates these updated preferences to the consent banner, enabling it to generate a new consent record when the user revisits the website.

4.1.5 Consent States Within the Record

The consent object uses a specific structure to represent different states. When a user has given consent to a specific service, the record includes when consent was given and in which consent record:

```
"matomo": {
    given_at: "2024-01-15T10:30:00.000Z"
    given_in: "550e8400-e29b-41d4-a716-446655440000"
    consent_initiator: "CB"
}
```

When a user has withdrawn previously given consent, both the original consent information and the withdrawal are recorded:

```
"google-analytics": {
    given_at: "2024-01-15T10:30:00.000Z"
    given_in: "550e8400-e29b-41d4-a716-446655440000"
    withdrawn_at: "2024-02-20T14:45:00.000Z"
    withdrawn_in: "660e8400-e29b-41d4-a716-446655440001"
    consent_initiator: "CB"
}
```

When no consent has been given for a service, the record contains an empty object:

```
"youtube": {}
```

When an entire purpose has no consented services, it may also be represented as an empty object:

```
"customise-ads": {}
```

4.1.6 Simplified UI Format

Within the user interface, consent is represented in a simplified boolean format that is easier to work with:

```
{
    "service-improvement": {
        "matomo": true,
        "google-analytics": false
    },
    "additional-features": {
        "youtube": false
```

```
    },
    "customise-ads": {}
}
```

The system automatically transforms between this simplified format and the full record format when loading and saving consent.

4.2 Consent storage

Consent records are stored across multiple locations, each with information density calibrated to specific functional needs. This distributed approach upholds the principle of data minimization.

4.2.1 Consent Store

The controller or a consent management platform (CMP) on behalf of the controller, stores the consent records (including withdrawals) for the service provider in a **consent store**. This serves as proof that the service provider is authorized to collect and process personal data based on valid consent. Such evidence is particularly important if the service provider must demonstrate compliance in the context of a complaint or regulatory procedure.

4.2.2 Subscribing Real-Time Events for Third Party Scripts

The consent record must also be made technically available to the controller in real time, enabling automatic activation and deactivation of third-party scripts. In other words, the controller and any TPPs it may use may only collect personal data (or, in the case of an opt-out process, must stop collecting personal data) if consent (or an objection) has been given. To achieve this, the controller must **subscribe** to the relevant events, such as granted or withdrawn consents, for each third-party service it intends to include in its service.

4.2.3 Client-Side Storage

A **cookie** is stored on the end user's device (e.g. in the browser) to ensure that previously made decisions are recognized during subsequent visits.

4.2.4 Consent Agent (History)

Each consent record is stored in the **consent history** of the consent agent, enabling users to review and modify their past consents (see 2.1.2).

The locally stored consent history contains a complete record, as it is used by the data subject to modify prior consents (for example, by withdrawing consent) and to demonstrate potential non-compliance by a controller. For a discussion of the decision not to encrypt this storage, see 6.3.

5. Consent mechanism (opt-in)

5.1 Legal pre-conditions

According to Articles 6(1)(a), 7, and 25(1) GDPR and Recital 32, consent must be a freely given, specific, informed, and unambiguous indication of the data subject's wishes, expressed by a clear

affirmative act. These requirements do not need to occur simultaneously; they may unfold as part of a sequence of user actions, provided that all conditions are fulfilled by the time consent is transmitted to the service provider (see 1.3). The following sections specify how these conditions are met by describing the user journey beginning with pre-settings in the consent agent over the handover notice to the cookie banner and contextual consent (and later on, a so-called privacy dashboard enabling data subjects to easily exercise their data subject rights).

Consent decisions are generally issued to a specific controller and tied to individual processing purposes but may later be adjusted with regard to specific data recipients. Where a controller engages processors or joint controllers, users who have previously consented to certain processing purposes may choose to opt-out of processing by individual data recipients—or opt back in at a later time.

The specification also defines rules on the continued validity of existing consents when relevant changes occur in the processing operations or purposes. It further specifies when and how controllers must re-obtain consent, which is no longer valid, and under which circumstances controllers may request consent again even if the user has previously refused (see 8.).

5.2 Consent Agent Pre-settings

The consent process is designed to begin within the consent agent, well before the data subject comes into contact with a specific controller or service.

5.2.1 UX/UI

The core function of the consent agent's user interface is to provide users with information about the processing purposes and operations, including their associated risks and benefits for the users' fundamental rights, for which service providers typically ask them for their consent. Based on this information, users can make first informed decisions regarding each purpose.

5.2.2.1 Informed

This requires the following information to be presented to the user in the consent agent:

- The **specific purposes** of the processing operations for which the user's consent will be transmitted to the controller if the controller requests consent for this purpose.
- The **categories of personal data**, which are typically collected for each purpose.
- The **categories of controllers, processors or other data recipients** that will typically process the user's personal data on the basis of their consent.
- Whether personal data might be transferred to third countries (non EU/EEA countries) and information on possible risks particular to data transfers to third countries without an adequacy decision and without appropriate safeguards pursuant to Art. 46 GDPR.
- Information on **the specific benefits and risks** to fundamental rights of the data subject resulting from the data processing operations for each respective purpose.

> **Example 5**
> The data subject is informed that, based on their consent for a given processing purpose, website operators process their data and typically engage third-party service providers for this purpose, such as statistically analyzing user behavior on the website or creating advertising profiles.

- The **legal basis** of the processing operation (e.g. consent according to Art. 6 (1)(a) GDPR or legitimate interests of the controller according to Art. 6 (1)(f) GDPR).
- The (at least approximate) **duration** the personal data is typically being stored and processed for.[22]
- If applicable, the **right of the data subject to withdraw** their consent (e.g. according to Art. 7 (3) GDPR).
- If applicable, information indicating that the data subject has the genuine **freedom to decide** whether to provide consent or not.
- Information on the **consequences** of giving or not giving consent.

> **Example 6**
> The data subject is informed that, when refusing consent to website analytics they may use the website as before, but the website is not able to improve their service for its users.
>
> **Example 7**
> The data subject is informed that when refusing consent, the website content is not accessible at all (only admissible in case of a paid alternative - “pay or okay”).

In addition to delivering the required information, users SHALL be provided with clear, easily understandable explanations that offer further context, tailored to the needs of an average user, on the following points:

- **Explanation of Legal or Technical Terms:** Provide clear, plain-language definitions for any legal or technical terms that may be unfamiliar to the average user—such as Controller, Joint Controller, Processor, Third Party, Vendor, Purpose, and Personal Data—to ensure all users can fully understand their meaning and relevance.
- **Description of Processing Operations Throughout the Data Lifecycle:** Present a straightforward overview of the data processing activities carried out for a given purpose. This should cover every stage of the personal data lifecycle, from initial collection to final deletion or irreversible anonymization.
- **Explanation of Risks, Benefits, and Stakeholder Interests:** Clearly outline how these data processing activities might pose risks or offer benefits to individual users and other stakeholders. Additionally, describe the specific interests and objectives that the controller and other involved parties have in performing these processing operations.

The required information SHALL be shown in a clear and simple language so it is understandable for the average data subject by avoiding overly complex legal and technical terminology[23] and be provided in an accessible form.[24]

The display of information SHALL be designed in a way that increases comprehensibility for the data subject. This includes use of established design principles and standardised icons (Art. 12 (7) GDPR).[25]

The information SHALL be displayed in an easily visible, intelligible and clearly legible manner and provide a meaningful overview of the intended processing.

---

[22] CJEU, 1. Oct. 2019, C-673/17, Planet 49, 72-79.
[23] See EDPB, Guidelines 05/2020 on consent under Regulation 2016/679 Version 1.1, 67; CNIL, Guideline on Cookies and Trackers 2020-091, 22.
[24] See EDPB, Guidelines 05/2020 on consent under Regulation 2016/679 Version 1.1, 67.
[25] See Art. 29 Data Protection Working Party, Guidelines on transparency under Regulation 2016/679, WP260 rev.01,12, 13.

5.2.1.2 Specific

Pre-settings SHALL be acquired separately for each processing purpose.

> The consent decision is specified with respect to the particular controller and processing operation during the handover stage (see 5.3).

5.2.1.3 Freely

The consent agent SHALL provide the data subject with a genuine and free choice by presenting the options to set pre-settings for giving and refusing consent for each purpose in an equally accessible manner, ensuring that refusing consent to the processing of personal data for a specific purpose does not require more effort than granting consent.

For this purpose the pre-settings SHALL be set to a **neutral ground position** by default.

5.2.2 Whitelisting of trusted controllers

The consent agent MAY provide users with the option to define purpose-specific pre-settings for a pre-determined group of (trusted) controllers. In such cases, the identity of all individual controllers to whom consent may be granted MUST be clearly and transparently disclosed to the user at the time the pre-settings are configured.

User decisions regarding such groups of controllers MUST be treated equivalently to other pre-settings. Accordingly, a consent decision in the form of a consent record is generated for each individual controller only when the user first visits or uses that controller's service. Upon this initial interaction, a handover notice MUST be presented to the user in the same manner as for standard pre-settings (see Section 5.3).

> The whitelisting approach may be supported by consent agent providers to enable users to grant consent to single trusted controllers and/or a defined group of trusted controllers. This can help increase consent rates for controllers that, on the one hand, enjoy a particularly high level of trust among their users and, on the other hand, rely on user consent to sustain their business models.
>
> This is especially relevant for the journalistic media and publishing sector, which often offers content free of charge and depends heavily on revenue from personalised advertising. Users of these services may therefore be willing to grant consent to selected publishers, for example to support independent journalism within the EU.
>
> The criteria for determining whether a controller is "trustworthy" in this sense are not addressed in this specification and are, for the time being, left to the respective consent agent providers to define and assess. However, efforts to identify and publicly disclose particularly trustworthy controllers are encouraged, as they can create meaningful incentives for higher data protection standards among businesses. This whitelisting approach may also be particularly relevant for controllers that have obtained certification of their data protection practices, for example through a GDPR seal pursuant to Article 42 GDPR.

5.2.3 Signals

The consent agent users' pre-settings are stored in the extension's persistent storage provided by the user's browser (see 5.4.4).

5.3 Handover

Upon first use of the controller's service (e.g., visiting the website), the consent agent initiates a handover process, in which the user confirms or adjusts their general pre-settings for that specific controller. If the user confirms or adjusts pre-settings so that they give consent for one or more purposes, a consent record is created for the specific controller.

Automatic handover of pre-settings
The user's consent agent presets are automatically transmitted to the controller after a defined waiting period. During this waiting period, the user may interrupt the automatic process by interacting with the handover notice or opening the consent banner. The user can then adjust their default agent settings before they are transmitted as specific consent to the service provider. Once the waiting period expires, any adjusted consent is automatically granted via the handover notice, or the user may manually provide consent through the consent banner.

5.3.1 UX/UI

A **Handover Notice** SHALL be presented to the user the first time they use a service after making any modifications to their pre-settings in the consent agent related to the purposes transmitted to that specific service provider.

If the user has modified settings in their consent agent, the handover notice reappears to confirm that these changes should be applied to the specific service (for example, a website).

The **Handover Notice** MUST include the explicit individual purposes or standardised and tested icons representing these purposes, as referenced in section 5.2.2.1.

The **Handover Notice** SHALL enable the user to

- modify their pre-settings previously configured in the consent agent specifically for the controller in question.
- open the consent banner of the specific service provider to obtain additional information on the details of data processing as required by section ….

As long as the end user is not making any modifications of their pre-settings in the CA, the **Handover Notice** SHALL only be displayed upon subsequent use of a specific service if there is a relevant change in the risks and/or benefits associated with a purpose, specifically when:

- The risk for a specific consent previously granted by the user has increased, or a new risk has emerged with respect to the specific service provider. In such cases, the service provider is obligated to **notify** the user or request **re-consent** in accordance with the requirements outlined in section 4.4 (Versioning).
- The risk has significantly decreased, thereby allowing service providers to request that users revisit their prior consent decisions for a given purpose—even if the user previously denied consent for that purpose with the specific service provider ("again-consent").

There are two reasons for permitting the service provider to request consent again after a user has previously denied it ("again-consent")—provided that the associated risks have significantly decreased (for example, through the adoption of less invasive processing methods, collection of fewer data categories, limitation of data recipients, or avoidance of data transfers outside the EU):

Most importantly, the risks and benefits of a processing purpose form the decisive basis for end users' decisions to give or withhold consent. If the risk-benefit ratio changes significantly, the end

user must be informed of this so that they can make a new decision accordingly.

> **Note:** If the risks increase or decrease significantly, the controller must inform users accordingly by triggering an appropriate notification (see 8.2).

Secondly, the possibility of asking the user again provides a positive incentive for service providers to make their systems increasingly privacy-friendly, especially if new privacy-enhancing technologies become available on the market.

Since risk reduction must be decisive in order to be able to ask for consent again, the frequency of requests should be kept within limits. Furthermore, the provider of the consent agent should monitor the frequency of requests and, if necessary, adjust the thresholds (which is a 'decisive' improvement).

The Handover Notice SHALL NOT reappear on different components of a service, such as subpages within a website.

The Handover Notice MAY automatically close after a sufficient[26] time interval that allows the user adequate opportunity to review their decisions regarding the specific service.

The handover notice ensures users can adapt their pre-settings for each individual controller and website they visit. This way all requirements set forth in recital 32 of the GDPR— clear affirmative action, freely given, specific, informed, and unambiguous consent—are met at the moment of consent.

The self closing mechanism of the handover notice ensures that the problem of "consent fatigue" is solved and that data subjects do not need to interact with a consent management UI anytime they visit a new website.

5.3.2 Signals (opt-in)

5.3.2.1 Message Passing Overview

The different components of the consent system communicate through message passing. This allows a compatible consent banner on a website to exchange information with the extension, even though they run in isolated contexts.

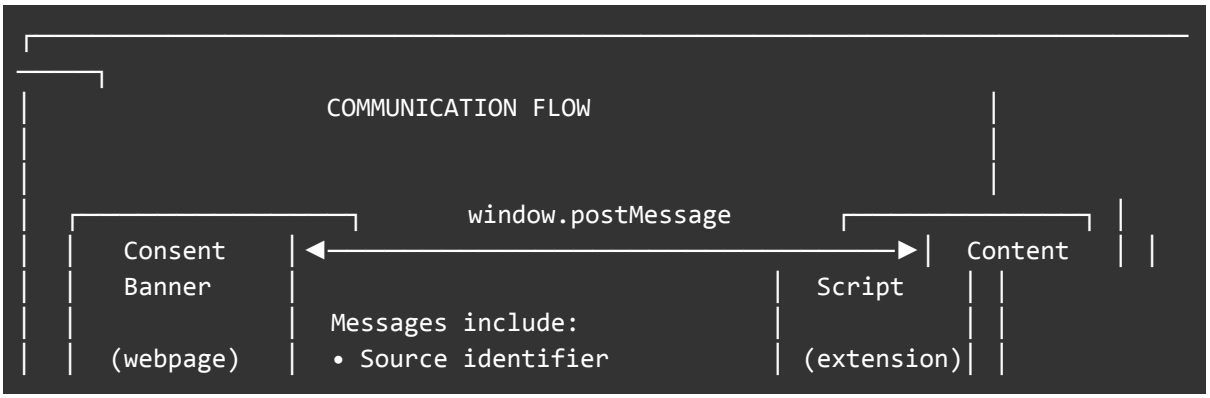

[26] Reference to the research apparatus and the past and upcoming studies performed therein to determine what is tested to be sufficient.

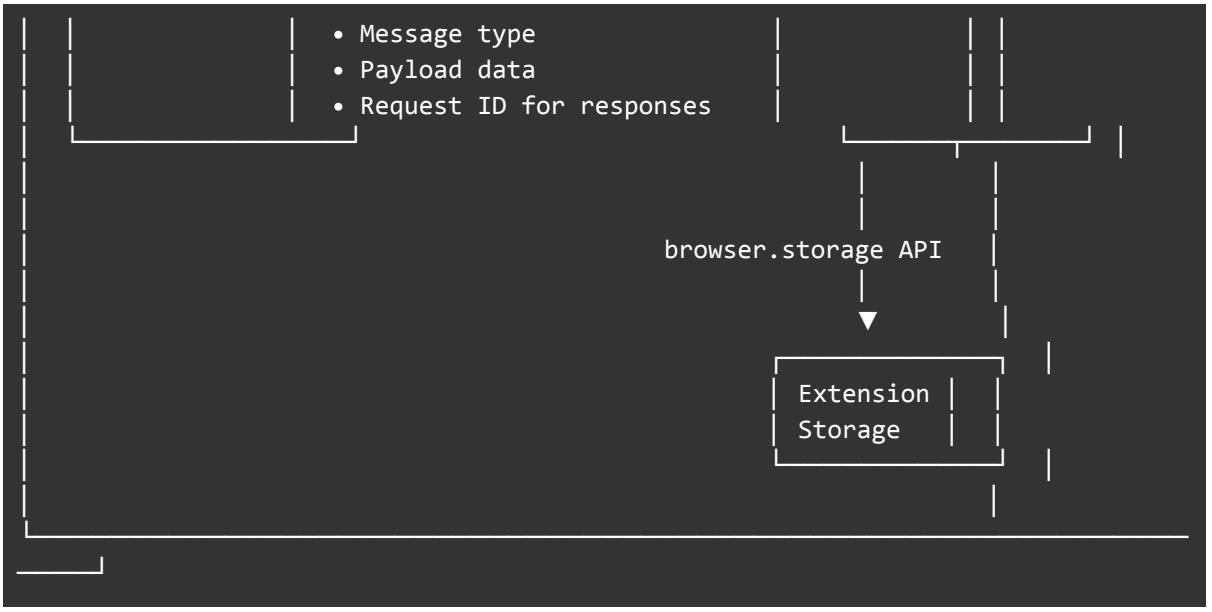

All messages include a source identifier that allows each component to verify where the message came from. The content script only processes messages from the consent banner and ignores all other messages on the page. This prevents malicious scripts from impersonating the consent banner.

5.3.2.2 Extension Detection

Before attempting to communicate with the extension, the consent banner must determine whether the extension is installed. This is done through a handshake process.

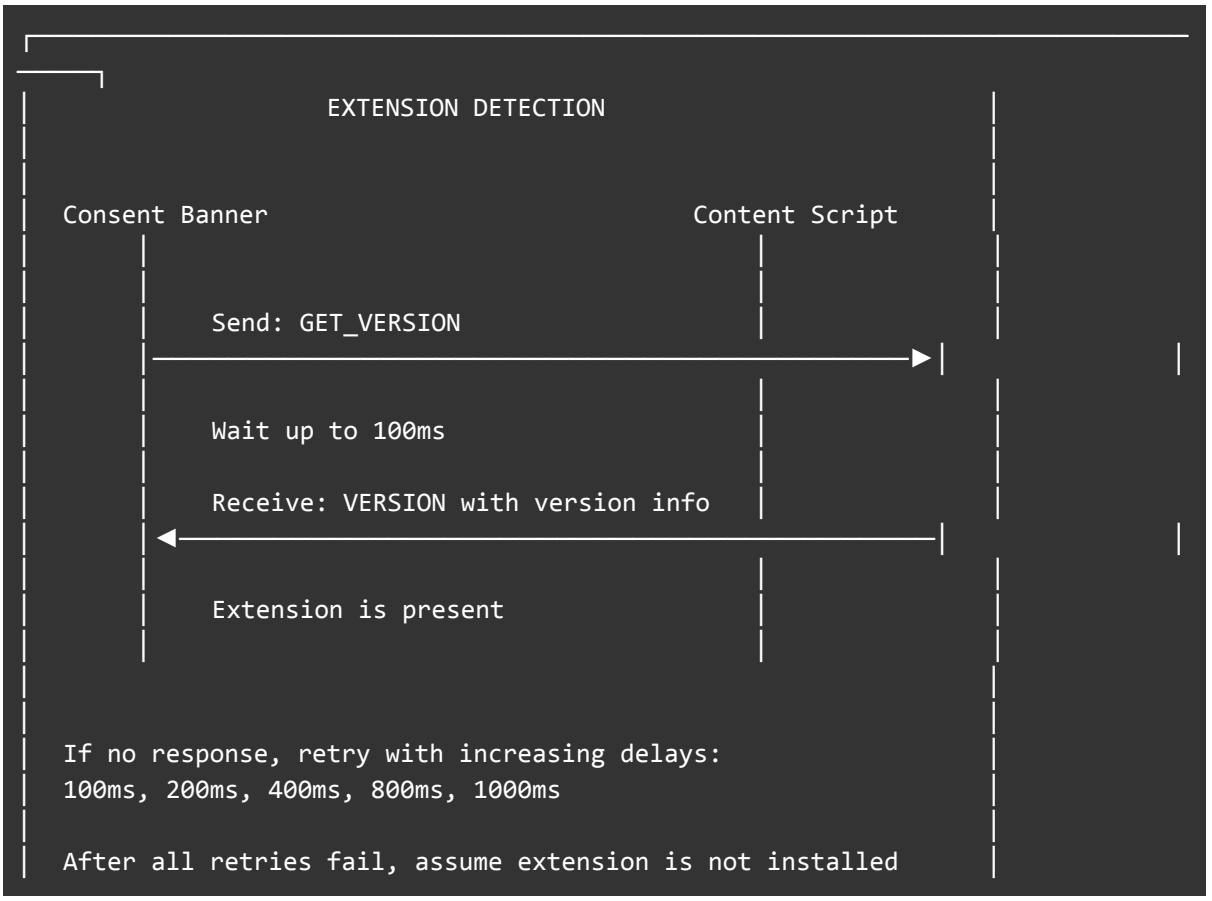

The retry mechanism with increasing delays accounts for situations where the extension might take time to initialize. Once detection completes, the result is cached so subsequent operations do not need to repeat the handshake.

5.3.2.3 Retrieving Consent Data from Extension

When a user visits a website and the consent banner determines that a consent agent is installed, the consent banner requests any existing consent data from the extension through the content script.

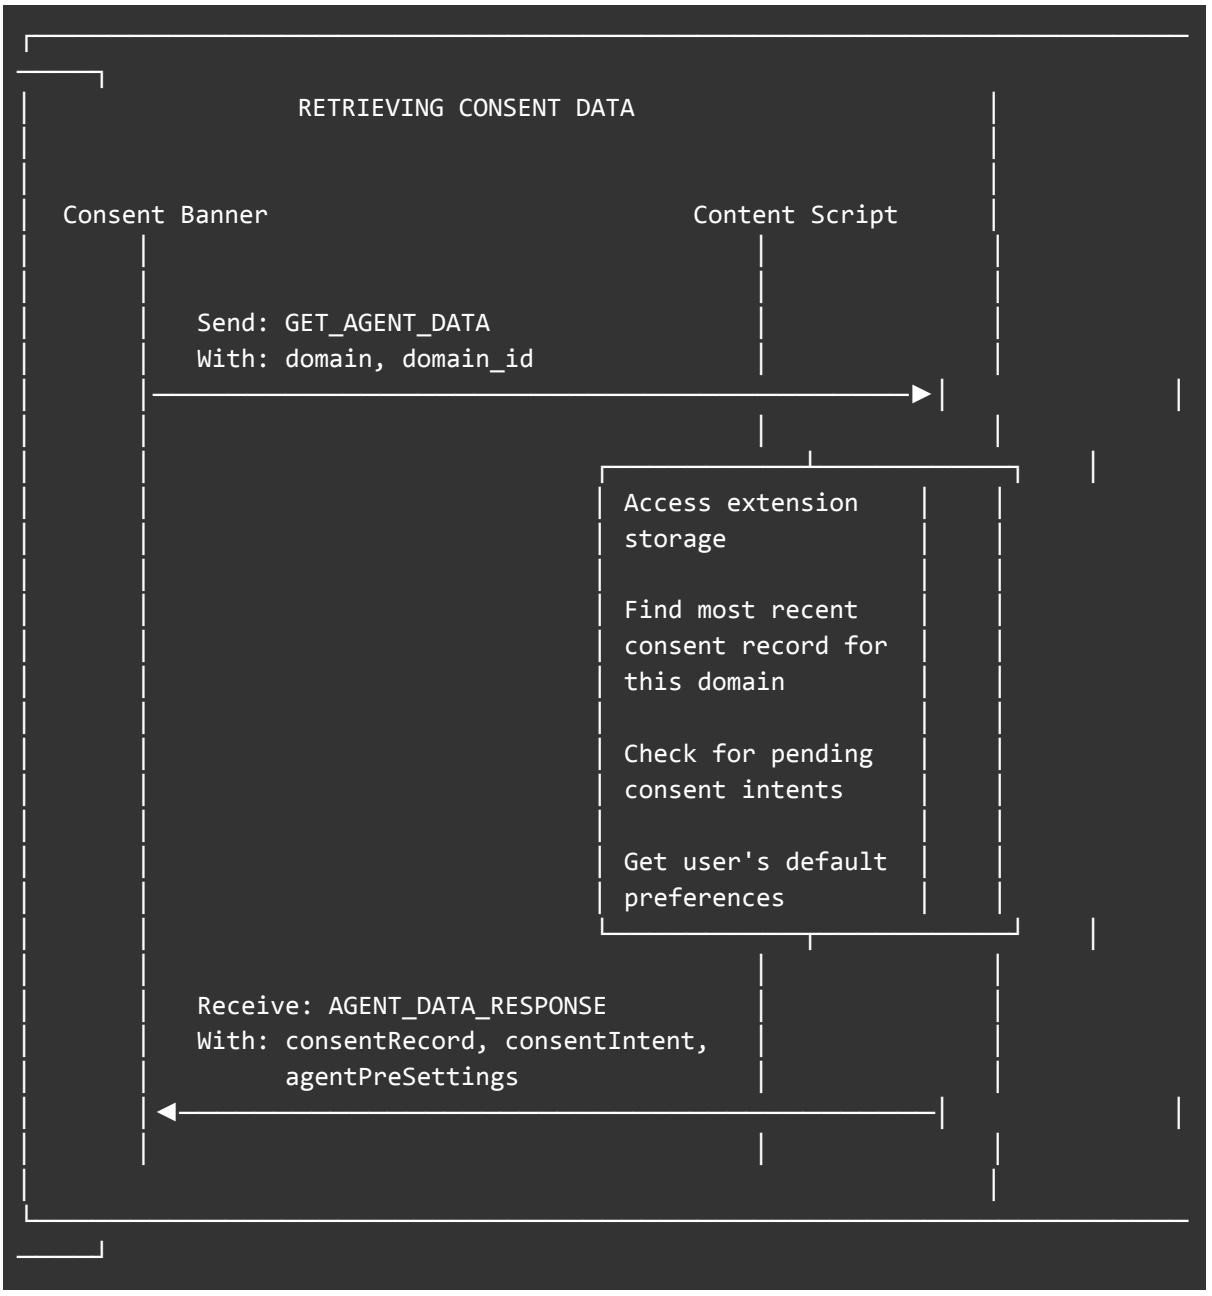

The response includes the most recent consent record for this domain if one exists, any pending consent intent the user configured before visiting, and the user's default preferences that can serve as starting points.

5.3.2.4 Saving Consent to Extension

When a user makes a consent decision and the server has processed it, the record is sent to the extension for storage through the content script.

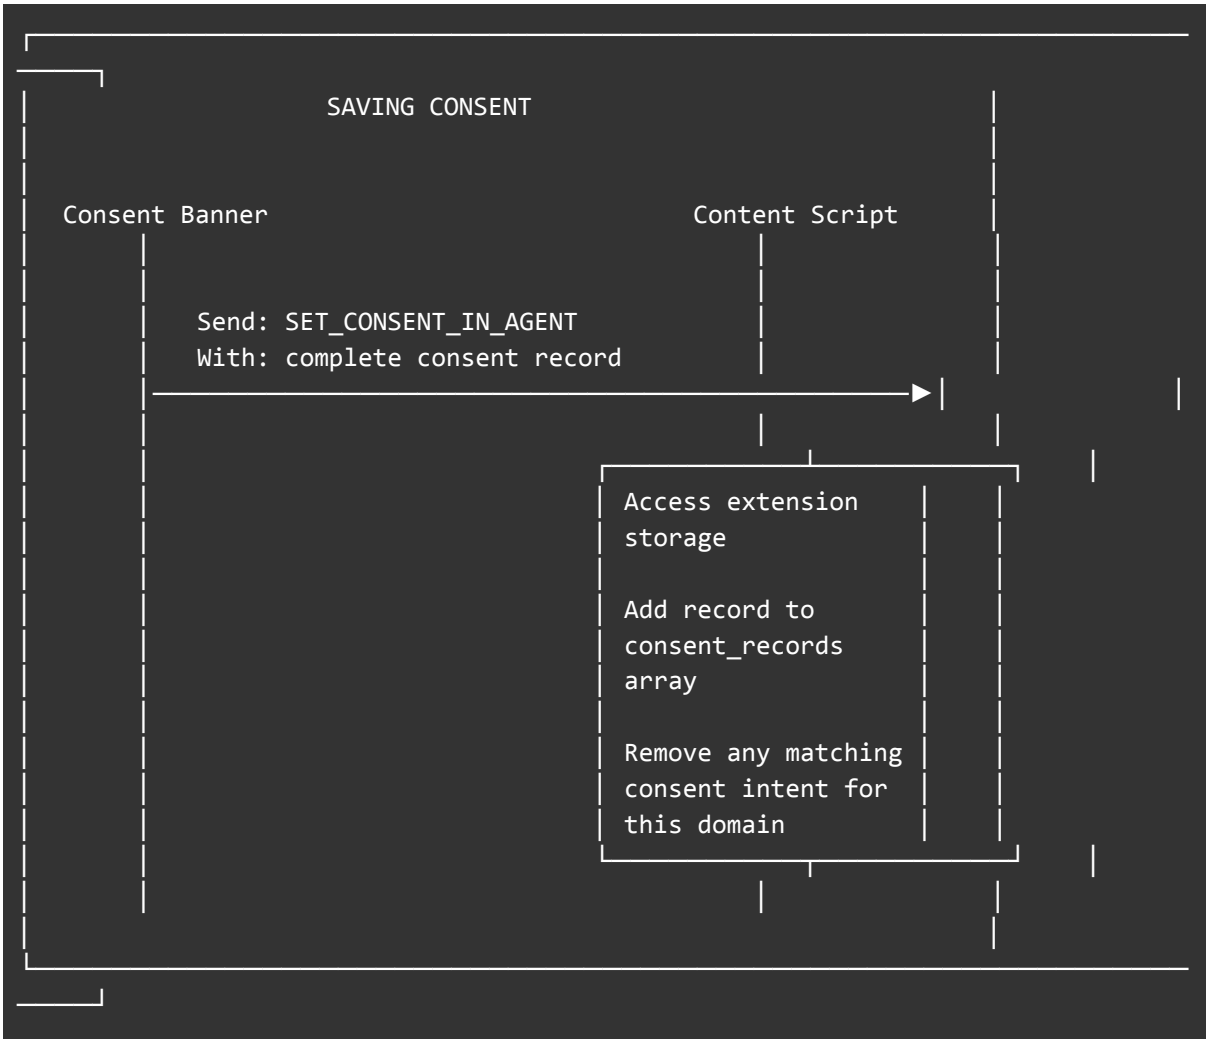

The content script appends the new record to the array of consent records in extension storage. If there was a pending consent intent for this domain, it is removed since the user has now made an actual consent decision.

5.3.2.5 Consent Saving Flow

For saving consent choices, the system follows a specific sequence to ensure the record is properly stored on the server before being saved locally in the different locations (i.e. cookie and consent agent).

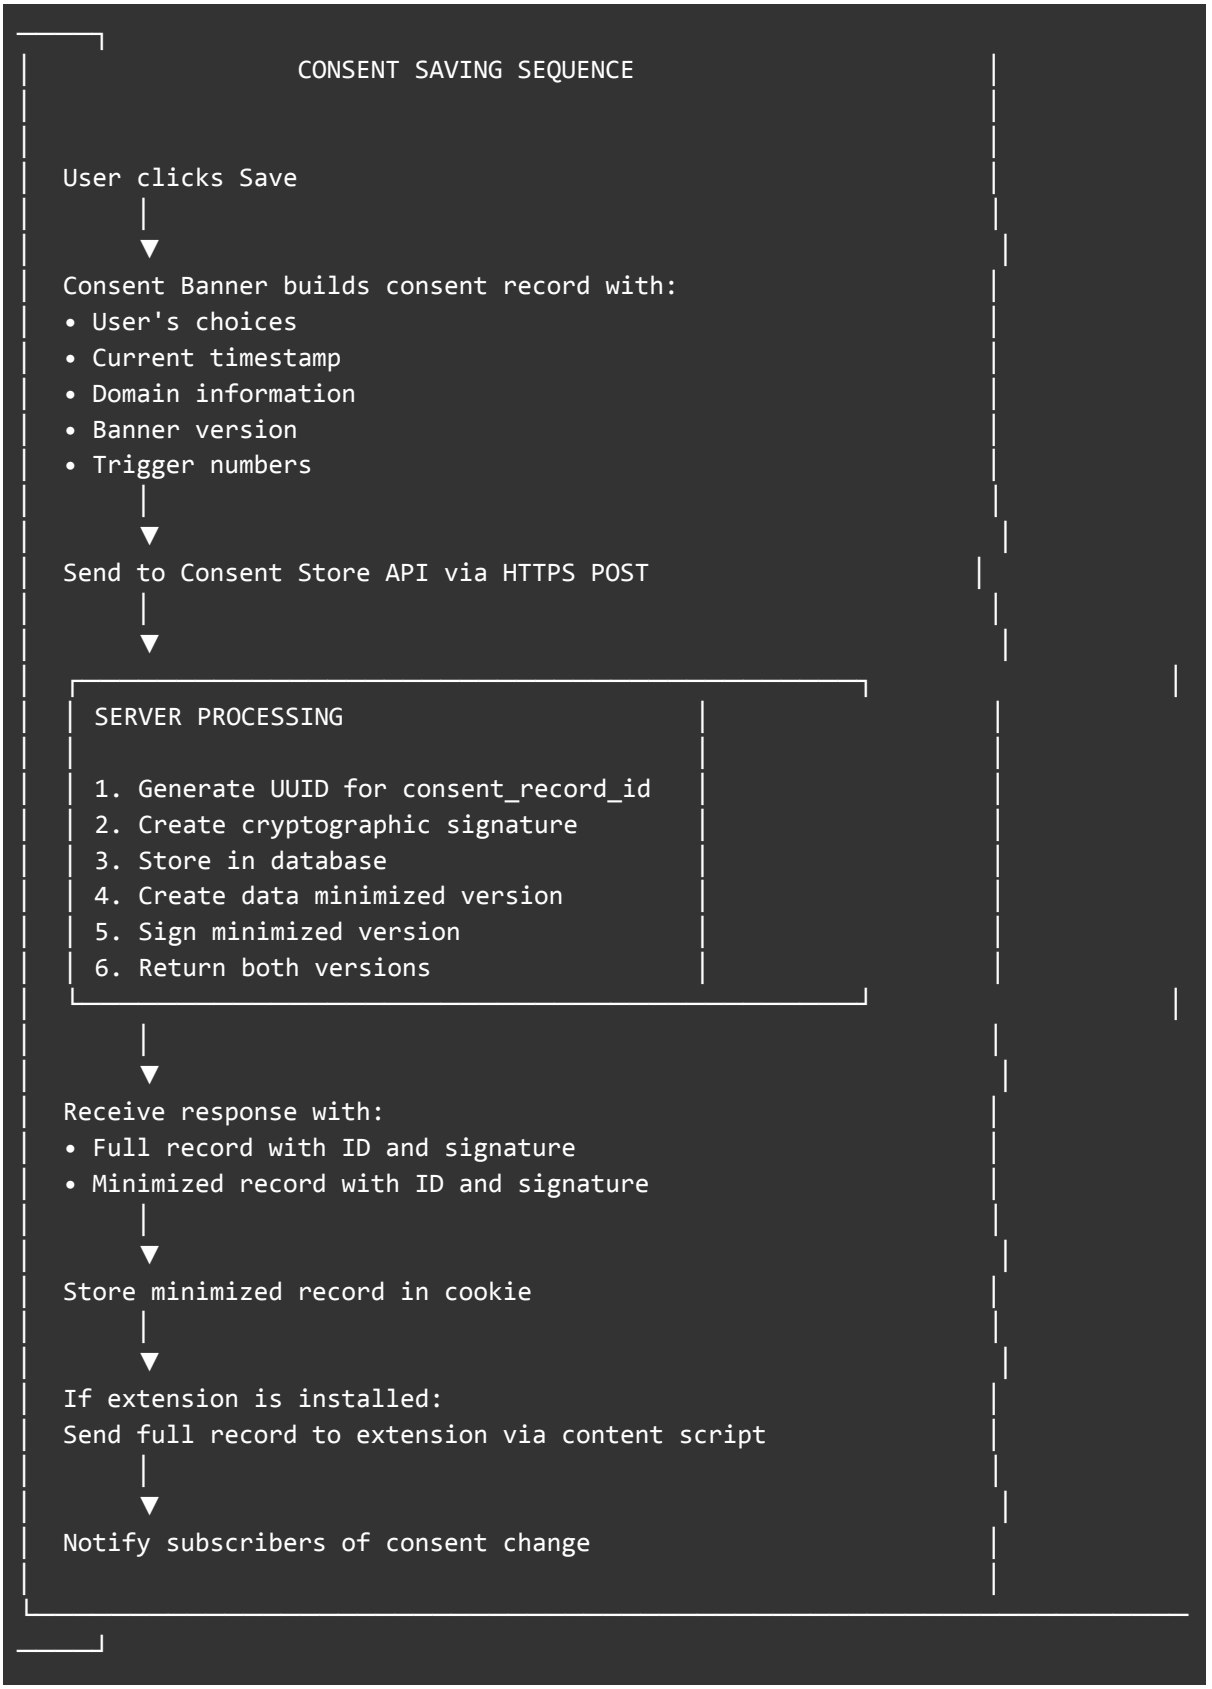

This sequence ensures that the server always has a copy of the consent record before it is stored locally. If the server request fails, the consent is not saved and the user is informed of the error.

⚠️ **There is one exception to this flow:** If the user declines all purposes (i.e., no consent is granted to any service), the record is stored locally without any communication with the server. Since the controller lacks a legal basis to process personal data based on consent in any form, there is neither a legal nor a technical justification (e.g., for evidentiary purposes or to facilitate third-party processing) for storing this decision beyond the user's device (see 5.4.2.1).

5.4 Consent Banner Interactions

The consent banner manages users' consent decisions for the controller. It handles both technical communication with the consent agent and conveys additional information to data subjects during the consent process.

5.4.1 User Journeys and Scenarios

This section describes how consent records flow through the system in different situations.

5.4.1.1 First Time Visitor Without Extension

This is the most common scenario. A user visits a website for the first time and does not have the browser extension installed.

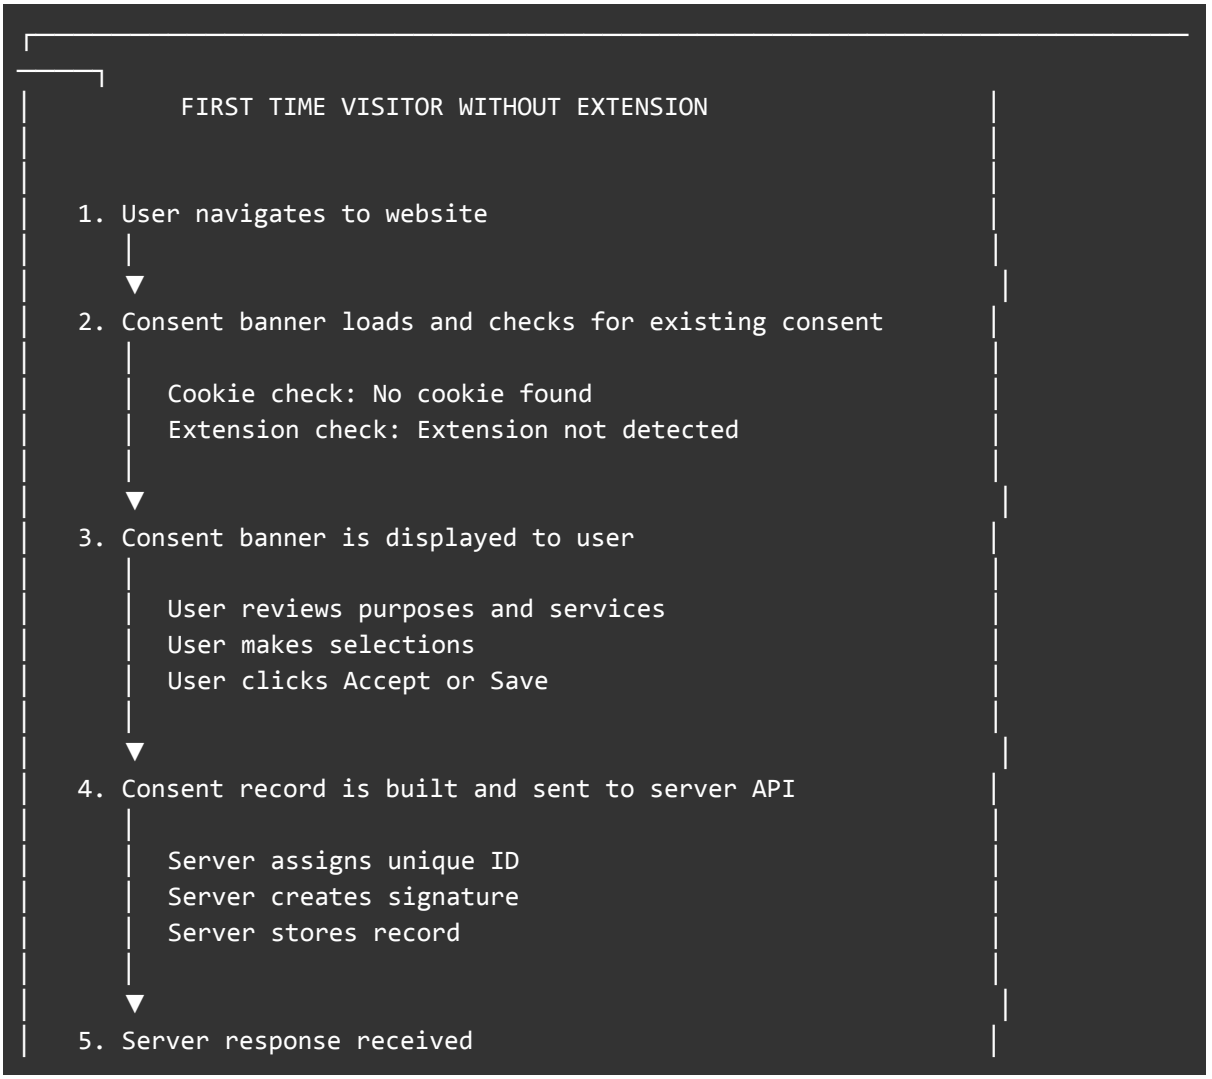

After this flow completes, the website can check consent status for any service and enable or disable functionality accordingly.

Call to action (CTA)
A broad adoption of consent agents by users requires that they are aware of emerging consent agents and their increasing privacy-enhancing capabilities.

A call to action should therefore be integrated into consent banners, directing users who do not yet use a consent agent to a common page (for example, a browser extension store) where all ConStand-compatible consent agents are listed. To enable this, the banner should, on first visit (when no cookie is set), detect that no consent agent is present.

If no consent agent is detected, the consent banner should display a CTA button—such as “Save for all sites”—within the banner when the user confirms their settings.

5.4.1.2 Returning Visitor Without Extension

When the same user returns to the website later, their previous consent is retrieved from the cookie.

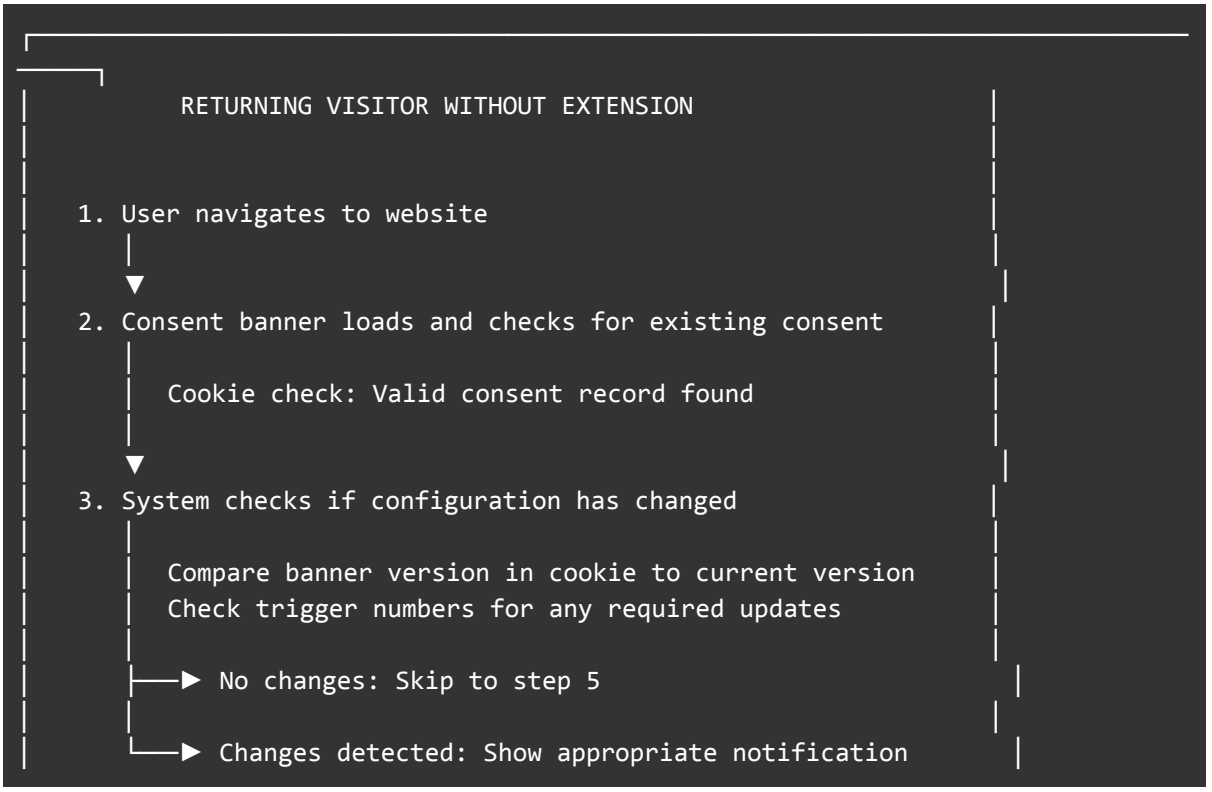

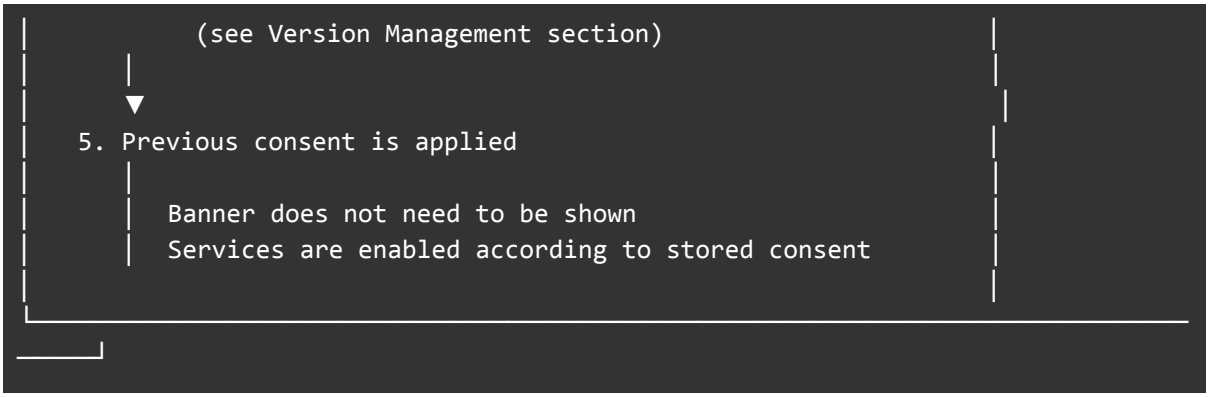

The key benefit of this flow is that returning visitors do not need to repeatedly make the same consent decisions. Their preferences are remembered and applied automatically.

5.4.1.3 First Time Visitor With Extension and Default Preferences

When a user has the extension installed and has configured default preferences, those preferences can be applied automatically with a brief confirmation period.

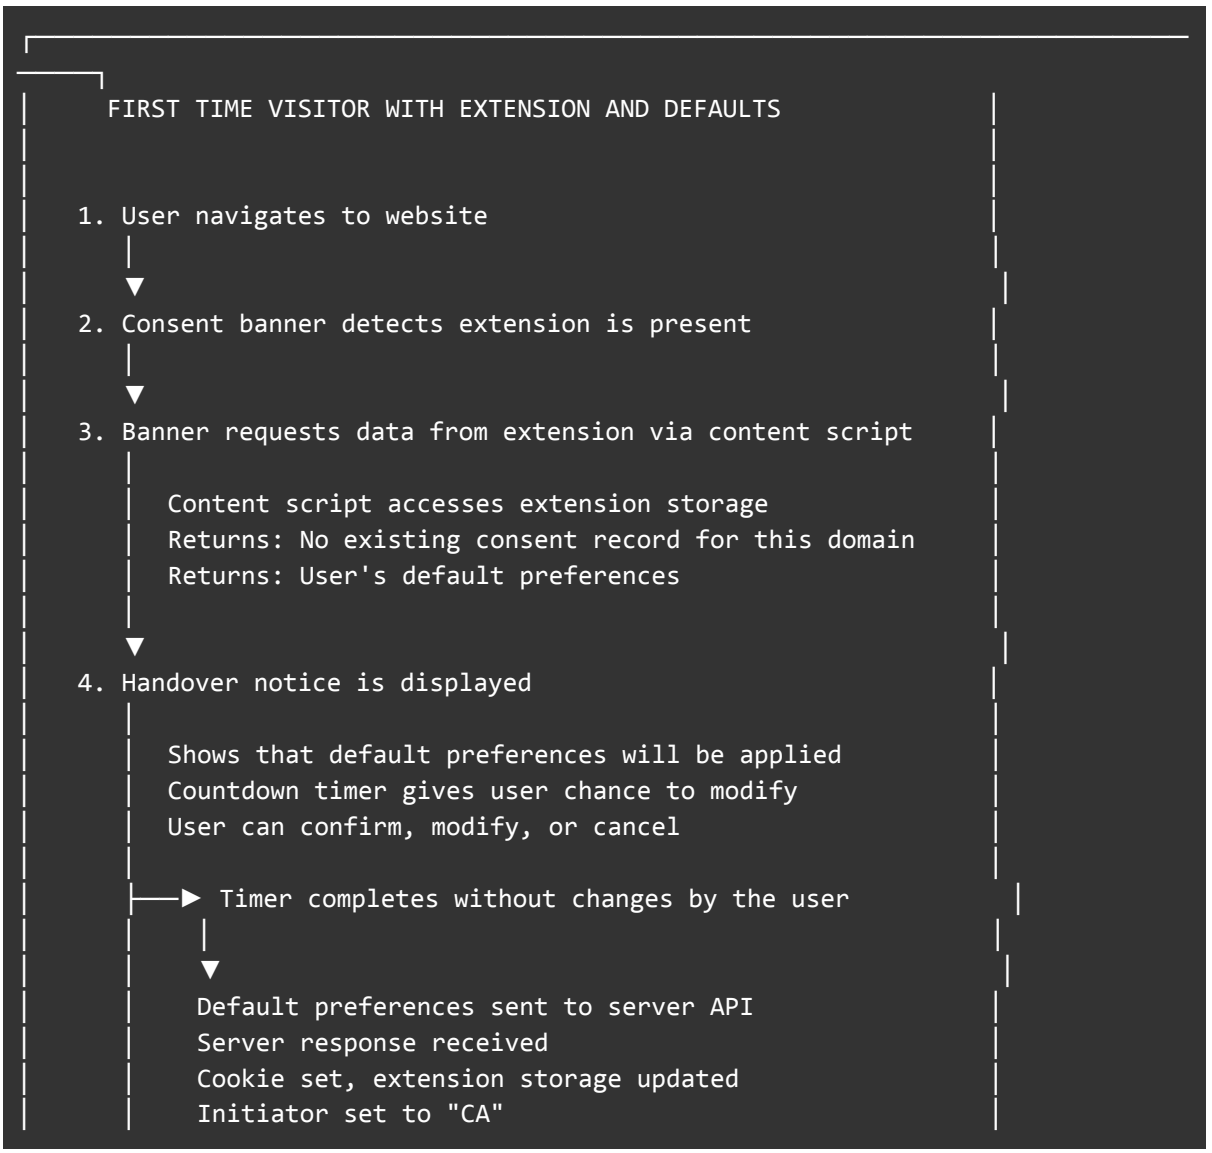

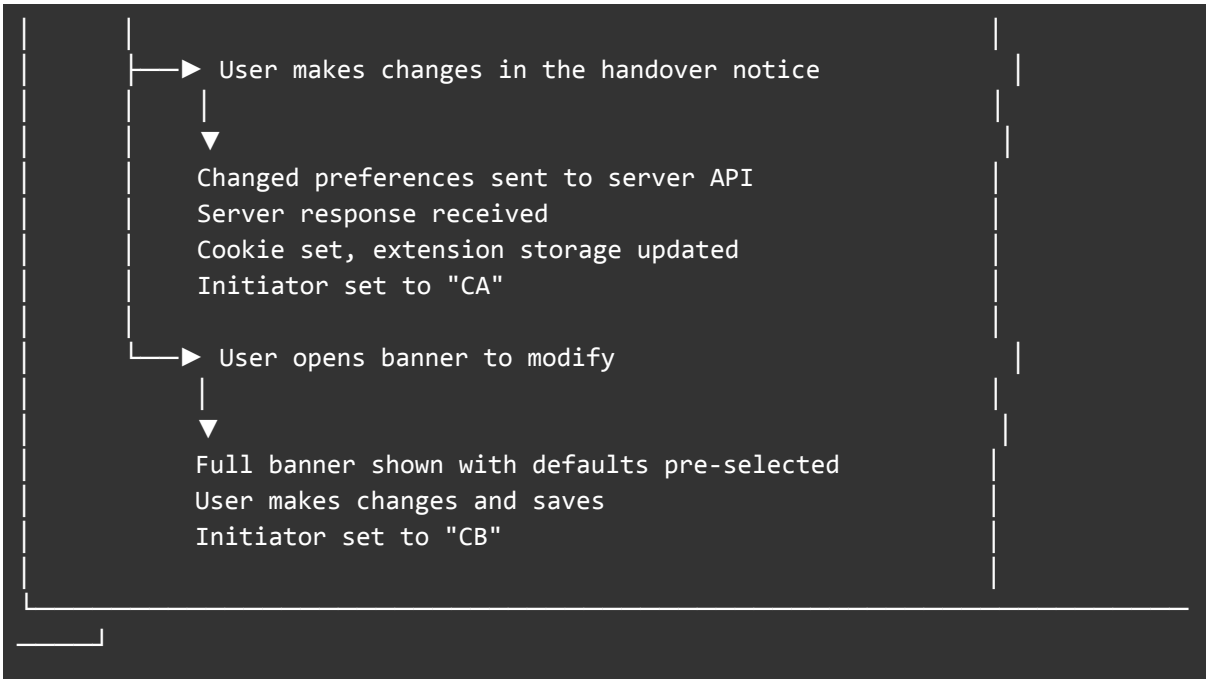

The handover notice provides transparency by showing users exactly what preferences are about to be applied, while still streamlining the experience for users who have already made their choices at the extension level.

It conveys the most important information, namely the purpose specific risks and benefits associated with the specific service.

5.4.1.4 Returning Visitor With Extension

When a user with the extension returns to a website, the system retrieves consent from both the cookie and the extension, then uses the most recent record.

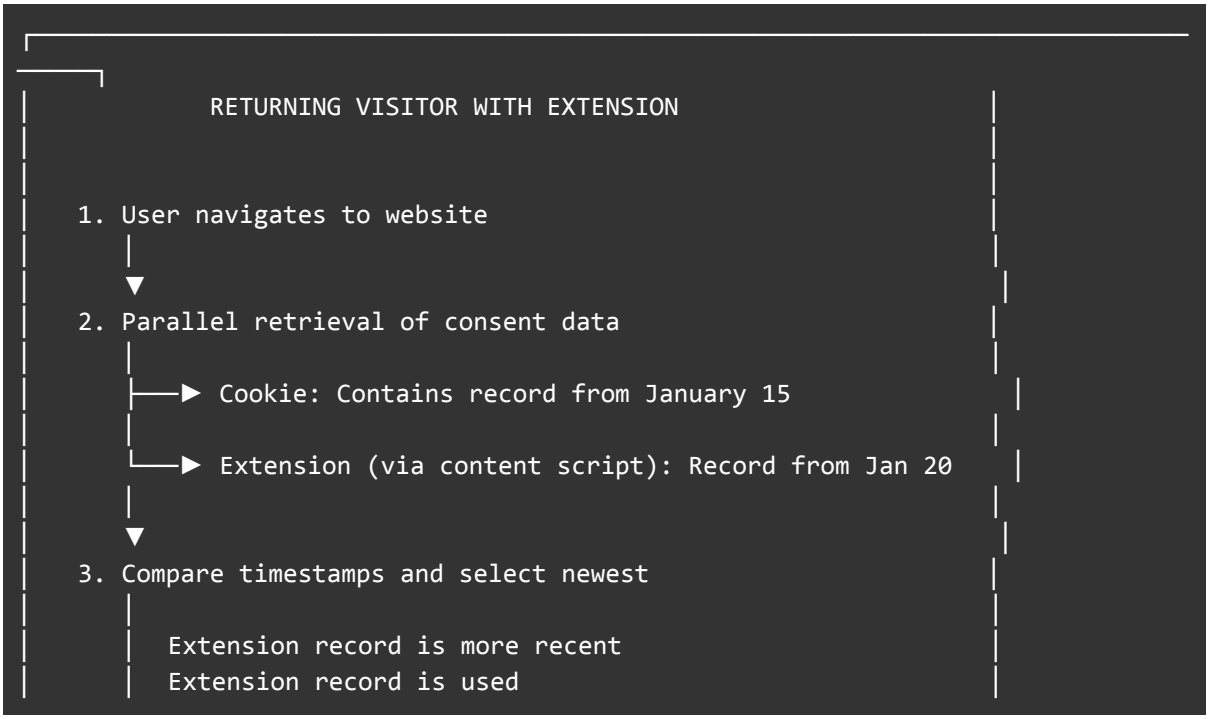

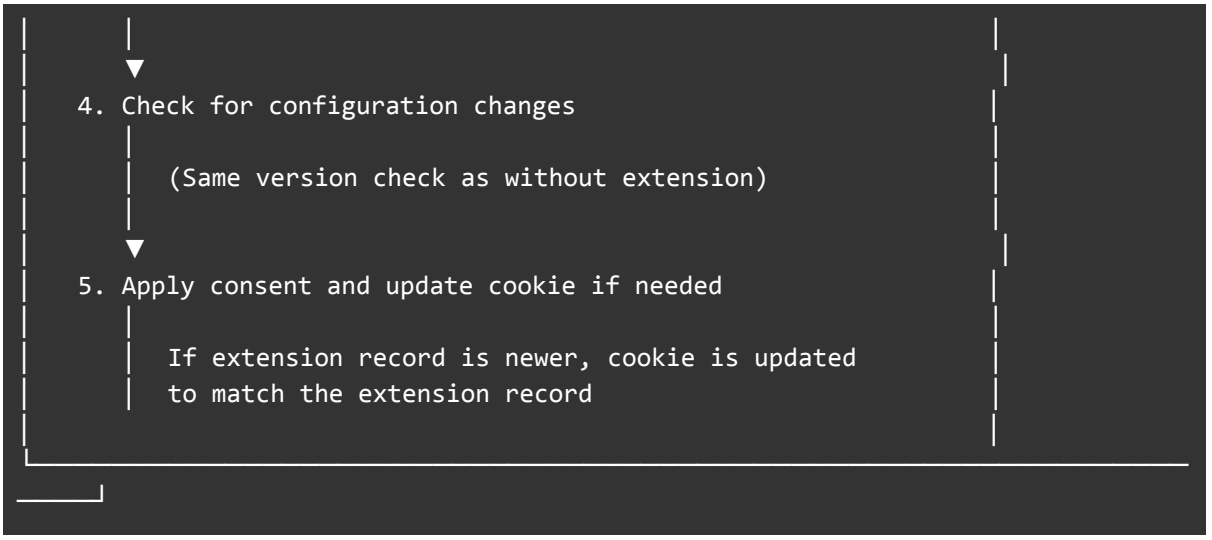

This ensures that changes made in one context (such as updating preferences through the extension interface) are reflected when the user visits the website again.

5.4.1.5 Contextual Consent for Embedded Content

If controllers embed third party content such as videos or maps in their service, they might want to collect consent contextually (e.g. within the respective embedded iframe) - see 2.3. Users can then grant consent in context when they want to interact with this content.

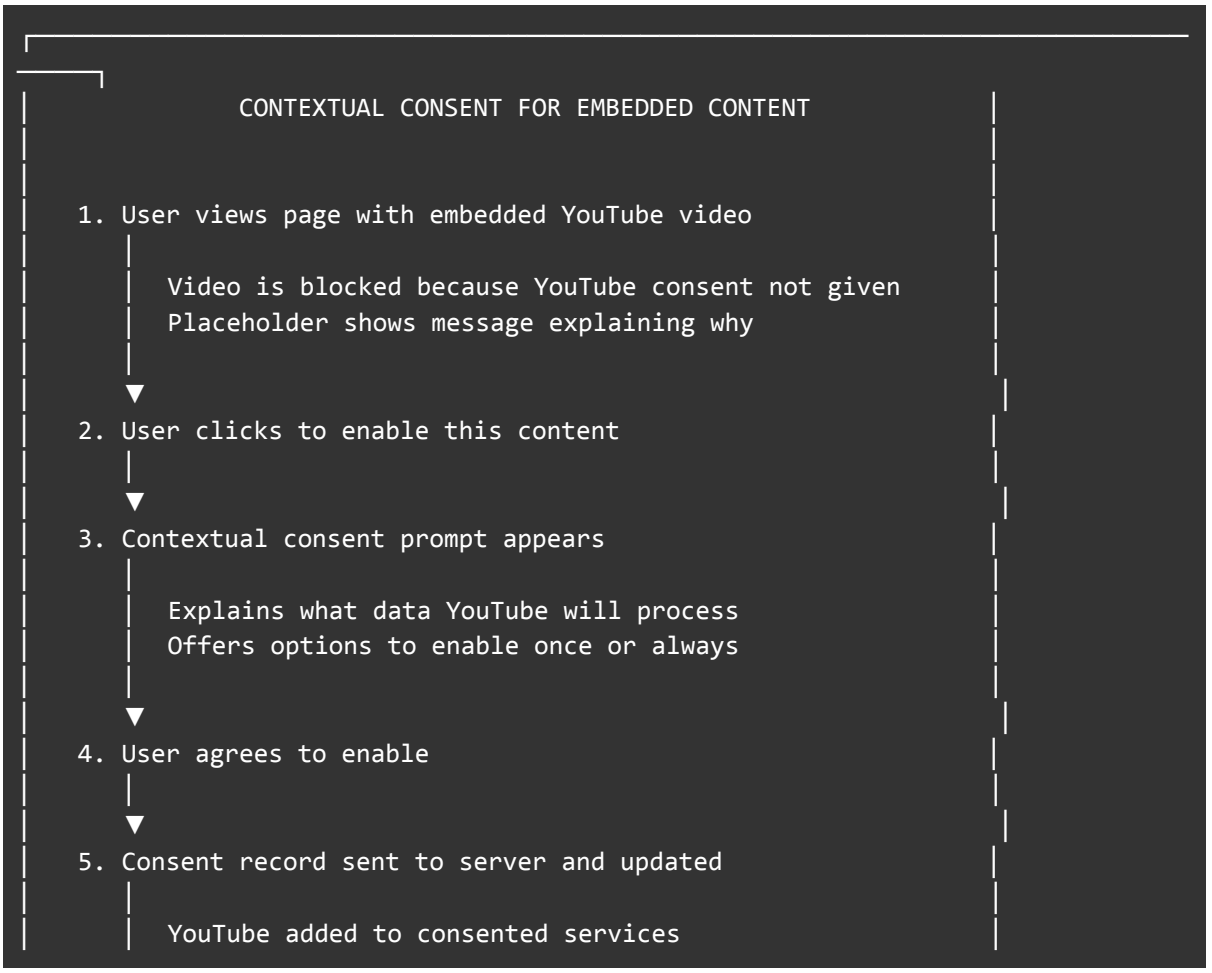

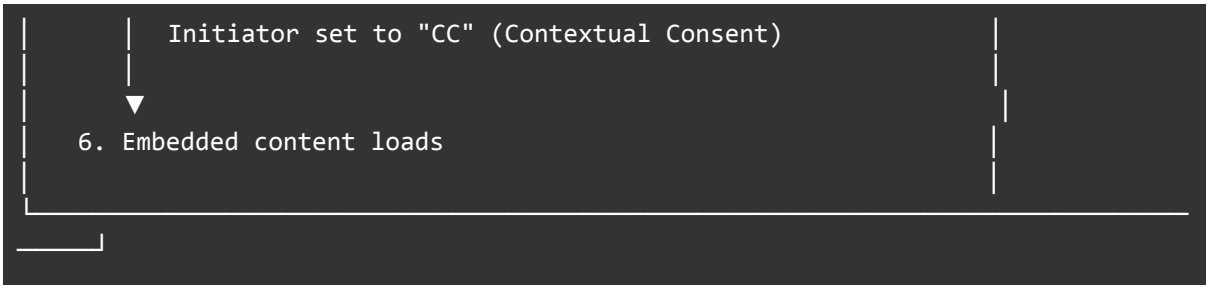

Contextual consent allows users to make just in time decisions about specific content without having to navigate back to the full consent interface. Contextual consent therefore enables users to better understand the significance of giving or refusing their consent (or of lodging their objection) in relation to this specific context.

5.4.1.6 Extension Not Responding

If the extension is installed but not functioning properly, the system falls back to cookie only mode.

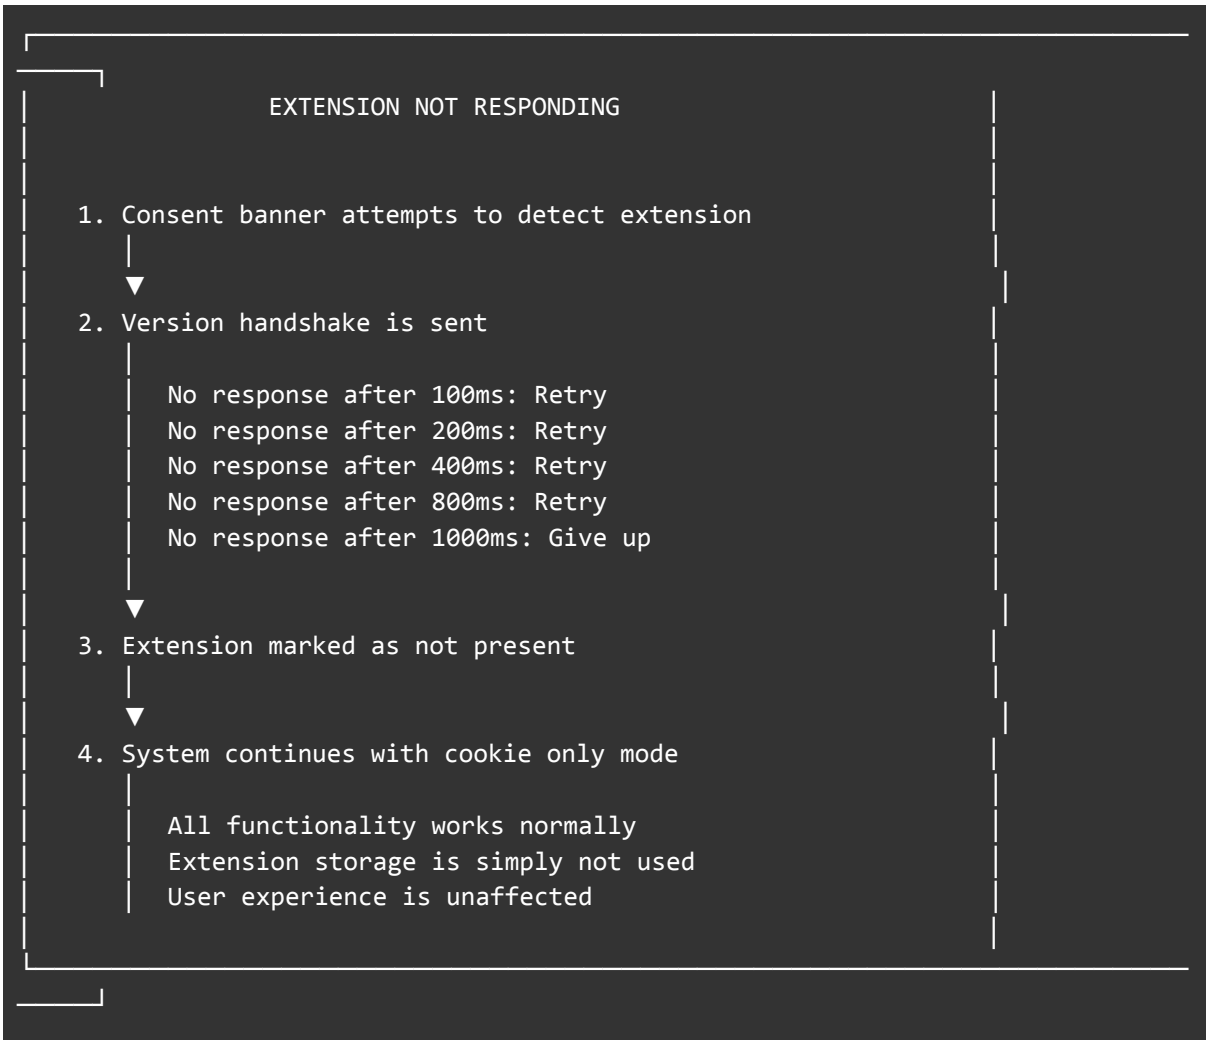

This graceful degradation ensures that technical issues with the extension never prevent users from managing their consent or using websites normally.

5.4.2 Configuration options

Data subjects can make different choices, causing different signals and storage operations. The system enables users to:

- Refuse consent (per purpose and per data recipient/vendor/third-party service provider)
- Give consent (per purpose and per data recipient/vendor/third-party service provider)
- Change previous decisions
- Confirm previous decisions
- Withdraw previous consent
- Renew consent after withdrawal
- Object to the processing based on legitimate interests (opt-out).

Users generally make these choices on **a per purpose basis.** The system is however set up in a way that stores consent on a **per data recipient basis**, enabling to de-select individual data recipients within one purpose, thus limiting the scope of their consent decision.

5.4.2.1 Refusing consent

Users may grant or refuse consent for each purpose individually.

Users **without a consent agent** can make their decision directly on a traditional consent banner, which appears automatically upon using the service (i.e. when visiting the website). They may consent to a purpose either for all data recipients/vendors/third-party services collectively or de-select individual services, when making their decision.

Users **employing a consent agent** may express consent through three distinct mechanisms:

(1) The general consent pre-settings are automatically transmitted from the consent agent to the specific controller via the handover notice (see 5.2), without any alteration by the user.

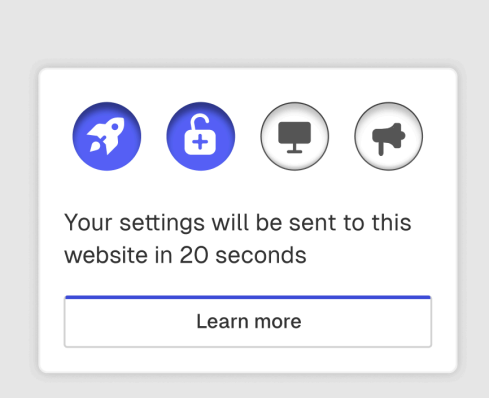

(2) The user alters their pre-settings in the handover notice for a specific controller by selecting or de-selecting purposes.

> **Example …**
> User U has declined consent for personalised advertising in the consent agent. When visiting the website of their favourite brand, U may choose to change this decision specifically for that controller C — for instance, to support the brand or to receive

personalised advertising based on their shopping activity on that site.

(3) The user opens the consent banner in order to make more granular choices regarding individual data recipients/third party services.

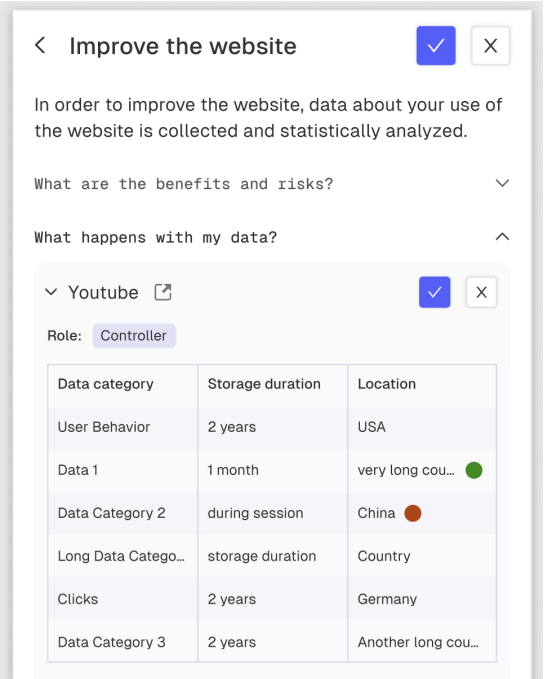

5.4.2.2 Changing or confirming consent

Users may change previous decisions.

If a user previously granted consent for a purpose (i.e., for specific services associated with that purpose) and later adds or removes individual services without revoking the entire purpose, this update is recorded as an amendment to the consent. Technically, consent already granted for one third-party service (Service A) is confirmed when the user later takes action (granting or revoking) for another service (Service B). The consent decision for Service A thus remains valid in the consent log, with the earlier grant referenced in the record.

5.4.2.3 Withdrawing consent

Users may withdraw consent at any time—either within their agent's consent history or via the service provider's consent banner.

According to the current state of the art, withdrawals made in the consent agent are retained until the user next visits the relevant service provider's service, at which point the withdrawal is forwarded to the provider as `consentIntent`. Legally, the withdrawal takes effect at that moment. The consent agent informs the user of the transmission through an on-screen notification.

As the state of the art evolves, withdrawals may be transmitted directly from the consent agent to the controller, allowing the withdrawal to take immediate legal effect. However, such real-time communication is not yet supported by the existing infrastructure.

5.4.2.4 Renewed consent after withdrawal

After withdrawing consent, users may grant new consent at any time.

5.4.2.5 Objecting (Opt-Out)

For processing purposes based on legitimate interests of the controller according to Article 6(1)(f) GDPR users may object to the processing of personal data (opt-out). Functionally, objecting to the processing works in the same way as withdrawing consent. The key difference is that the respective purpose is already pre-selected upon the user's first visit, enabling controllers to collect personal data **immediately** upon service use, as long as the user has not objected to the processing before (see 7.).

5.4.3 UX/UI

To support both consent agent users and non consent agent users, a consent banner must fulfill diverse UX/UI requirements. The banner must enable **consent agent users** to inform themselves about specific purposes and processing operations if desired, without obstructing service use. The banner must enable **non consent agent users** to make informed, specific, and free consent decisions.

5.4.3.1 Informed

When opened, the consent banner SHALL display all information referenced in 4.1.2.1, but tailored to the specific controller's processors and processing operations.

5.4.3.2 Specific

Consent SHALL be acquired separately for each processing purpose.

If consent is given for a purpose, all data recipients SHOULD be automatically activated by default.

Users SHALL be able to individually deselect data recipients within a single purpose.

5.4.3.3 Freely

For consent agent users, the consent banner MAY NOT open automatically upon the user's visit, but only upon their manual request.

For consent agent users the consent banner SHALL display, once manually opened by the user (either via the HN or the floating button - see 5.4.3.4) the consent decisions as received by the user via the handover mechanism.

> **Example 8**
> Consent agent user U visits controller C's website for the first time.
>
> U's pre-settings (affirmative for purposes 1+2; refusing purposes 3+4) are transmitted to C via the handover notice.
>
> C's consent banner remains closed until U manually opens it. Upon opening, the banner accurately reflects U's choices, with purposes 3+4 set to refused (not neutral).

For obtaining consent of non consent agent users the decision mechanism SHALL be set to "Off" by default or start off in a **neutral default state**, enabling the user to make an explicit decision for or against the processing of personal data.

For providing an “Opt-Out” mechanism[27] to consent agent users and non consent agent users the decision mechanism MAY be set to "On" by default in the consent banner.

The decision mechanism on the banner in any case SHALL NOT employ choice architecture techniques that manipulate users into providing consent (i.e. dark patterns). To this end:

- In a multilayered Consent Form, the options to **accept** and to **reject** the processing of personal data SHALL both be presented on the same layer, preferably on the initial layer. The option to reject consent SHALL require an equal or fewer number of clicks than the option to give consent.
- The options to **accept** and to **reject** the processing of personal data SHALL have matching text treatment and, for each, a minimum contrast ratio of 7:1.[28]
- The unselected **accept** and **reject** options SHALL in general have the same background color. The **accept** option MAY ONLY differ in color if it does not draw more attention than the reject option (for example, by being brighter or more inviting).[29]

If the banner is deployed automatically for non consent agent users and covers parts or all of the service’s content, thus impairing the accessibility and usability of the service, the banner SHALL be capable of being closed or moved out of sight by the user without giving consent (no “Cookie Wall” practice).[30] This requirement does not apply for appropriate “Pay-or-Okay” models.

If a paid subscription or similar fee is offered as the only alternative to the processing of personal data for a given purpose (“Pay-or-Okay”), such an alternative is permitted, if the **fee is appropriate** with regard to:

- The significance of the service for the data subject, for example, its relevance for participation in social life;
- Any potential position of power of the service provider in relation to the data subject; and
- The economic benefit derived from consent for the respective purpose.[31]

If the banner enables refusing consent simply by closing it or not interacting, this SHALL be clearly indicated to the user.[32]

If a non consent agent user sets preferences for all purposes in the banner, the banner MAY ONLY reopen automatically in cases of a significant risk increase or decrease - see section 5. Version management.

5.4.3.4 Persistent access to the banner

The consent banner SHALL be easily accessible with a single click from any part of the service (e.g. from every subpage of a website) for the user to adjust their settings. This MAY, for example, be implemented by means of a floating button or a link in the footer.

---

[27] For example in cases referred to in Article 21(1) GDPR.
[28] WCAG level AAA.
[29] See CNIL, Recommendations on Cookies and Trackers 2020-092, 34; DSK, OH Telemedien 2021, Version 1.1, 133,134; EDPB, Report of the work undertaken by the Cookie banner or other Consent Form Taskforce 2023, 18.
[30] See EDPB, Guidelines 05/2020 on consent under Regulation 2016/679 Version 1.1, 39-41; CNIL, Guideline on Cookies and Trackers 2020-091, 18.
[31] See ECJ, Case C-252/21, ECLI:EU:C:2023:537, 150.
[32] See CNIL, Recommendations on Cookies and Trackers 2020-092, 33; ConPolicy, Good Practice Initiative, Design Guideline 2023, page 7.

5.4.4 Signals and Consent Storage

This section explains each storage mechanism and when it is used.

5.4.4.1 Server Storage (Consent Store)

When a consent record is created, it is sent to the Consent Store API for permanent server side storage. This is always the first step before any local storage occurs.

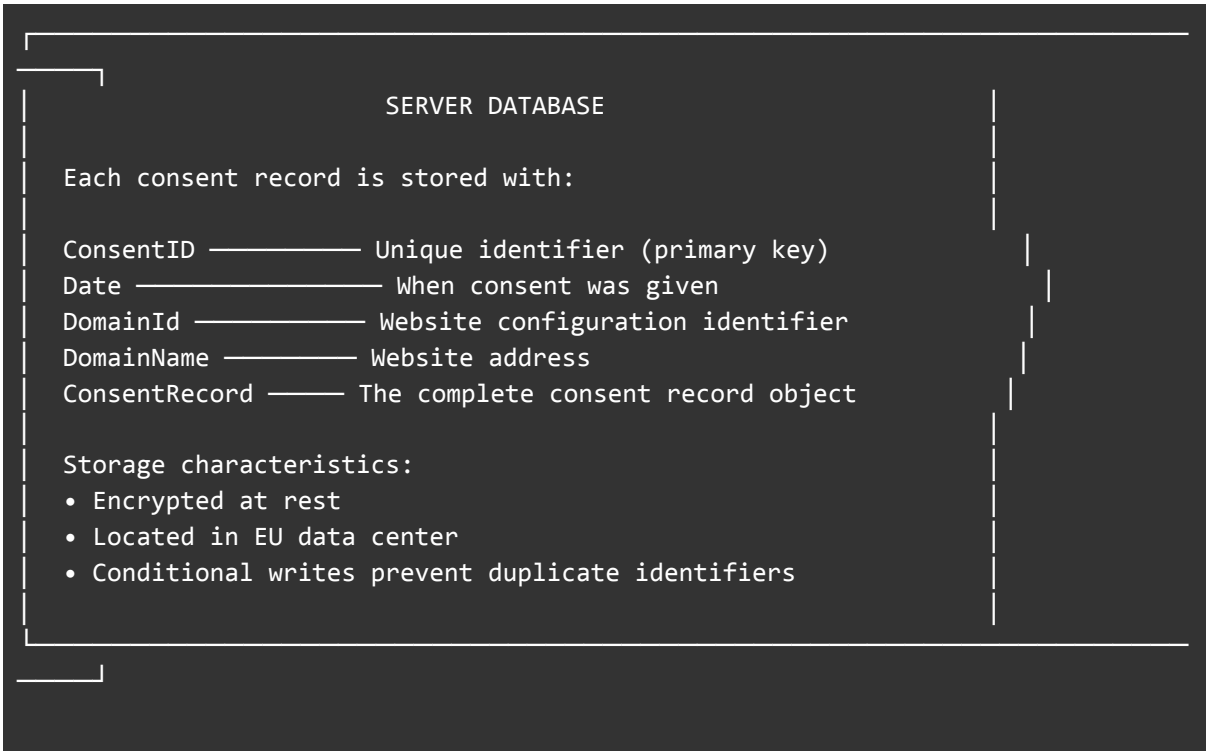

The server generates a unique identifier for each record using UUID version 4. Before storing, the server uses a conditional write that fails if the identifier already exists. In the statistically unlikely event of an identifier collision, the operation is retried with a new identifier.

The server also creates a cryptographic signature for each record. This signature uses an elliptic curve algorithm (ECDSA with SHA-384) and is created using a key stored in a hardware security module. The signature can later be used to verify that a consent record was actually created by the respective consent banner and has not been modified (e.g. by another malicious extension - see 6.4.1 Cryptographic Signature).

5.4.4.2 Cookie Storage

After a consent record has been processed by the server API and returned with its unique identifier and signature, the consent banner stores a copy in a browser cookie. This cookie serves as the primary local storage mechanism for users who do not have the browser extension installed.

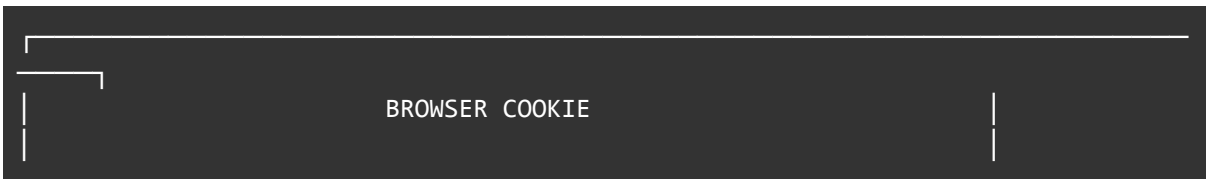

The cookie contains the data minimized version of the consent record, which includes all active consent information while omitting detailed withdrawal history. The strict same site policy prevents the cookie from being sent in cross origin requests, providing security against certain types of attacks. The secure flag ensures the cookie is only transmitted over encrypted HTTPS connections.

When a user returns to the website, the consent banner reads this cookie to determine if the user has already made a decision (e.g. on giving or refusing consent or objecting to the data processing purposes). If a valid record on a decision exists and no significant changes have been triggered by the controller (see 8.2), the banner SHALL not be shown again.

> Asking for consent in the consent banner after a user has refused to provide consent for a given purpose, is thus only possible, if the trigger system allows for it (see 8. Version management).
>
> Asking for consent contextually on the other hand remains possible also if a user has refused to provide consent before (see 2.5 Contextual consent).

It is important to note that the cookie is only set after the server has successfully processed the consent record. This ensures signal integrity, i.e. that every locally stored consent record corresponds to a record that exists on the server.

5.4.4.3 Extension Storage

When the browser extension is installed, consent records are also stored in the extension's dedicated storage area. This provides several advantages over cookie only storage.

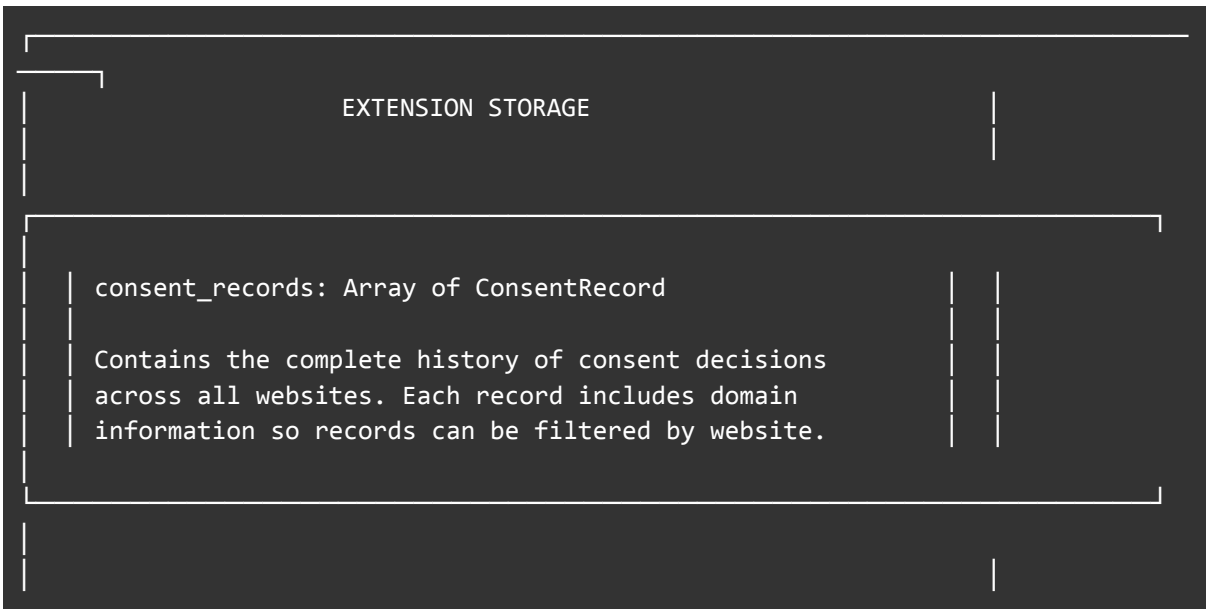

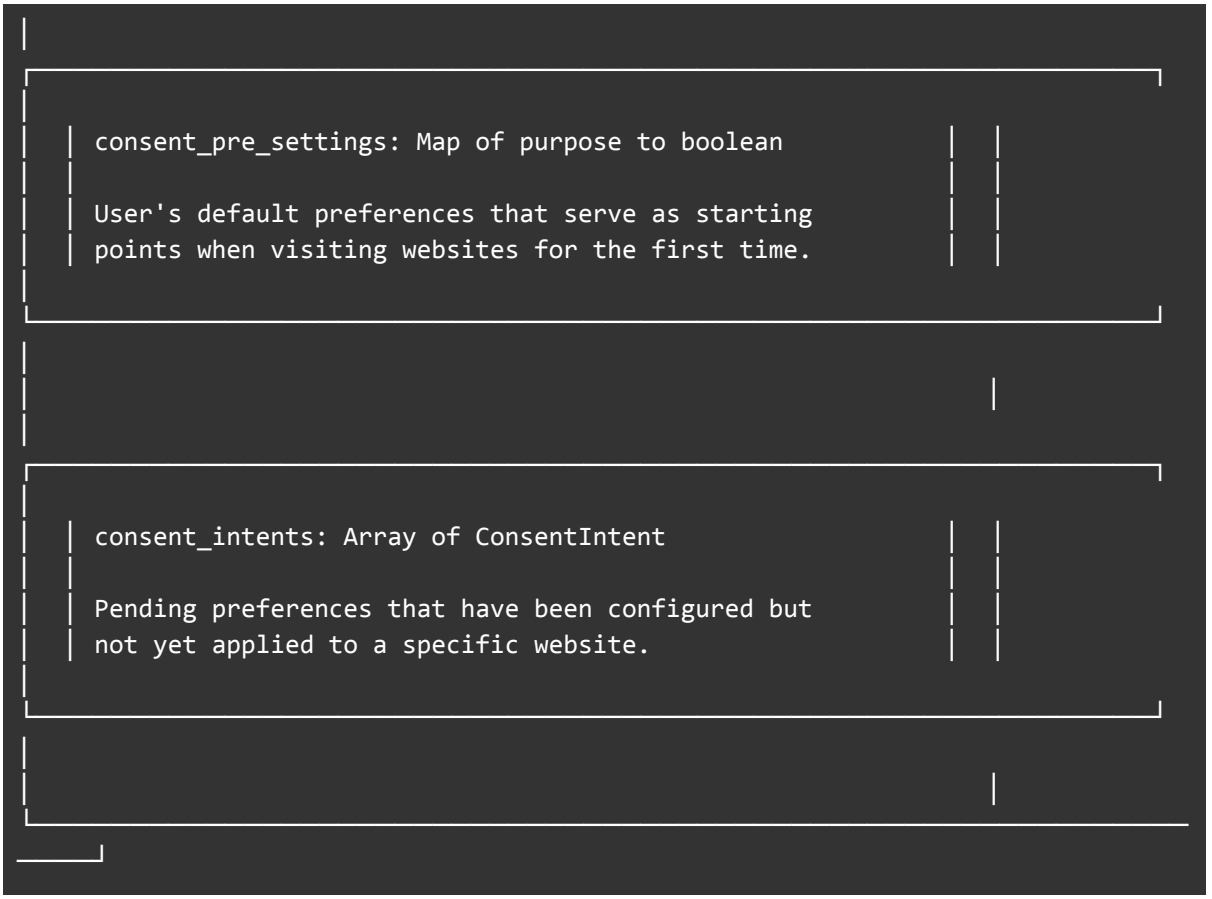

Extension storage is not bound by cookie limitations. It can store an unlimited number of records, is not cleared when users clear their cookies, and persists as long as the extension is installed. This makes it ideal for maintaining a comprehensive consent history.

The extension storage is local to the device. If a user has the extension installed on multiple devices, each device maintains its own separate storage.

Cross-device synchronisation
To improve both availability and user experience, user data MAY be synchronised across different devices. This COULD be achieved via a cloud service that aggregates and consolidates the data only temporarily and then returns it to the respective devices, without storing it permanently in the cloud.

“Consent Wallet”
Another possible development involves local data storage on the user’s device within an encrypted database of a native application. Such an architecture is currently attracting considerable attention, for example in the context of the EUDI Wallet initiative (see 6.4.2).

5.4.4.4 Storage Priority and Synchronization

When multiple storage locations contain consent records for the same domain, the system must decide which record to use. The decision is based on the timestamp in each record.

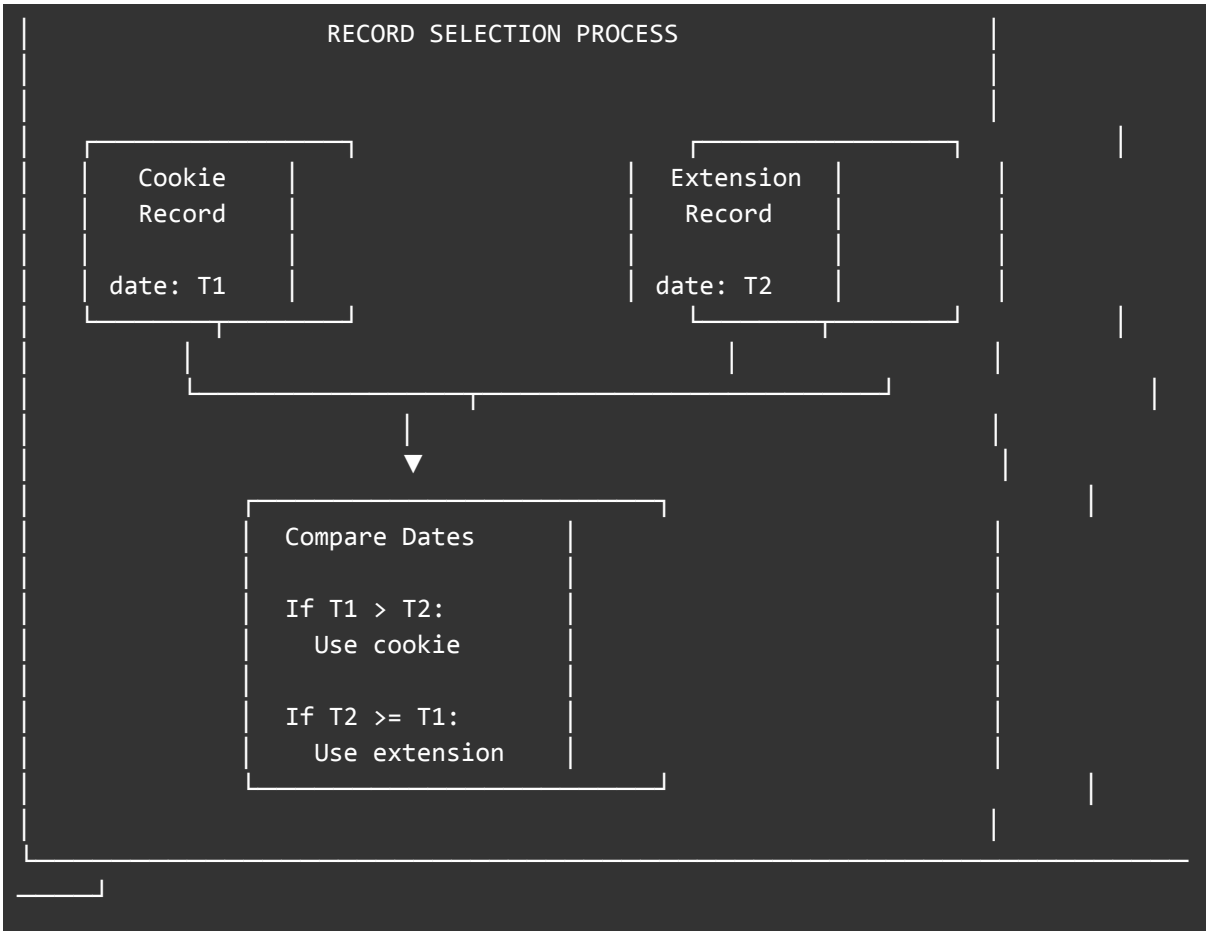

The most recent record is always preferred. This ensures that if a user updates their consent in one location, that update takes precedence over older records in other locations.

6. Data Protection and Security Considerations

6.1 Considerations on Data Minimization

The system adheres to the principle of data minimisation, meaning each **asset** contains only the information that is **functionally** necessary. Summarised from the previous chapters:

6.1.1 Functions

For the service provider

- Requests for consent (in opt-in procedures)
- Receipt of withdrawals of consent granted
- Receipt of objections (for opt-out procedures)
- In future versions: Receipt of data subject rights requests

For data subjects (users):

- Granting of consent (opt-in) or refusal
- Withdrawal of previous consent
- Adjustment of consent
- Submission of objections (opt-out)

- In future versions: Submission of data subject rights requests
- Avoidance of consent fatigue, i.e. no renewed requests unless circumstances have changed significantly
- Overview of all consents, withdrawals and other requests.

For both:
- Proof and verifiability in cases of dispute that consent was given (usually required by the service provider) or not given (usually required by the user).

6.1.2 Assets

Consent Pre-Settings: A set of binary choices (yes/no) indicating whether the user wishes to give consent for a specific purpose (or not), or to object to the processing for that purpose (or not).
It is therefore a combination of the user's settings in the agent.

> The pre-settings serve the purpose of generating a consent record in a consent banner (CB) on a website after the consent pre-setting has been transferred to the CB via Handover Notice – HN, while being manually adjustable in the website-specific context by the user.

Consent records: The documentation of the user's specific consent. The records consist of different elements/information, depending on the storage location:

- **Consent store** (see 5.4.4.1 and below), containing:
  - Purpose-specific consent decisions per third-party service provider (granted, confirmed, withdrawn, objected)
  - Metadata on consent decisions per third party provider or purpose (ID, time stamp, CB version, website domain). **NOT included** are for data minimisation purposes refused consents.
- **Cookie** (see 5.4.4.2), containing:
  - Contains purpose-specific consents per third-party service provider (granted, confirmed, withdrawn, objected **and** refused)
  - Metadata on consent decisions per third party provider or purpose (ID, time stamp, CB version, website domain). ID or timestamp are **NOT included**, if consent for all purposes has been refused, ensuring that users cannot be tracked by the service provider.

    > The **cookie** serves the purpose of automatically transferring and applying previously granted and refused consents when revisiting the service provider's website.
    >
    > It is:
    > - **Essential** for non-agent users, as there is no other way to (temporarily) store granted consents.
    > - **Useful** for new agent users, as the agents can reuse and continue to use the existing consent records from the cookie.
    > - **Useful** for agent users if the agent is temporarily unavailable.

- **Consent agent** (see 5.4.4.3), containing
  - Purpose-specific consents per third-party provider (granted, confirmed, withdrawn **and** refused)
  - Metadata on consent decisions per third party provider or purpose (ID, time stamp, CB version, website domain).

- Additional agent-specific characteristics, e.g. information on whether a withdrawal had already been transmitted to the service provider (e.g. when re-visiting the website).

  This **consent agent** storage serves the purpose of automatically transferring and applying given and refused consent decisions (e.g. when re-visiting the website).

  The storage of all decisions (including refused consents) in the agent enables multi-device/multi-browser use by the user, whereby their current decision is always taken into account.

(Partial) Storage of Refused Consents: Data Minimisation vs. Synchronisation
The use of a consent agent limited to a single browser or end device does not necessitate storing instances in which the user has refused consent. In such cases, the principle of data minimisation pursuant to Article 5(1)(c) GDPR requires that these refusals not be retained.

However, if a user wishes to apply the consent agent across multiple browsers or devices at a later stage, enhanced usability may require that refusals of consent be stored. Synchronising the consent history of the first consent agent with that of the subsequently installed one ensures that all previously defined user settings are transferred.

If refusals were neither stored nor synchronised, the user would again be prompted for consent on each website previously visited, despite having already expressed refusal.

As the implementation of such synchronisation is a concrete objective, the storage of refused consents for this purpose is not considered incompatible with the principle of data minimisation.

Consent History (Agent): A collection of consent records (agent) for a user, documenting the history of the user's consents. It

- is stored in the user's agent
- enables consent to be withdrawn via the agent
- may be used later to resolve disputes. For this reason, withdrawn consents remain stored.

Consent Store: A collection of consent records received by a service provider from different users. It

- documents the consents relating to the service used (e.g. website) given by users when using the service.
- is stored in the technical system of and managed by the service provider (or a trusted third party (e.g. a CMP) on behalf of the service provider).
- may be used later to resolve disputes. For this reason, revoked consents remain stored here as well.

6.1.3 To What Extent Do the Assets Contain Personal Data?

[...]

6.2 Limitation of data use to consent management

Considerations on limiting the data use to what is necessary for consent management:

Consent Pre-Setting

- Non-identifying and intentionally de facto public

- Technically minimal and tailored to functionality such that it is of little or no use beyond the intended purpose
- Transmitted only locally between the agent and the consent banner.

Consent Record (Consent Store)

- Transmitted in encrypted form in transit between the consent banner and the consent store.
- Stored in the consent store (must be appropriately secured there).
- Transmitted to the server-side system (Trusted Time Stamping Service) for signing (see 6.4).

Data is stored centrally and must be protected against unauthorized access through technical and organizational measures, thereby minimizing the risk of misuse (see 6.2.2).

Consent records are personal data as they document user consents, containing a unique Consent ID and creation timestamp.

Single Record (Repeated Submission)
Weakly identifying: Submitting the same record twice allows re-identification of a user across website visits via the Consent ID or timestamp, though with minimal privacy insights.

Linked Records (Updates)
Explicit predecessor ID enables user recognition to mark prior consents obsolete—a deliberate functional necessity for the service provider, remaining weakly identifying.

Unlinked Records (Different ID/Timestamp)
Potentially weakly identifying: More consent variations than pre-settings mean "exotic" record pairs may probabilistically link to the same user (k-anonymity considerations).

Aggregated Records
Combinations enabling stronger linkability occur only in the Consent Record Store; isolated records pose only weak re-identification risk, warranting protection.

Note: Despite the very limited linkability of multiple consent records pertaining to an individual, the service provider—or any trusted third party storing the data on their behalf—must implement state-of-the-art data security measures.

Consent Record (Cookie)

- Transmitted in encrypted form in transit between the consent banner and the service (e.g. website).
- The cookie is origin-bound, meaning that only the controller's service can access it.
- The cookie is not marked as HttpOnly, so it can be read via JavaScript, which is technically necessary for the service's functionality.
- However, JavaScript outside the controller's service cannot access the cookie. Only the service provider can access it.

Consent Record (Agent)

- Transmitted only locally between agents and consent banners.
- Stored in the Consent History (Agent) and appropriately secured there.
- Technically signed in the backend of the agent provider and contractually safeguarded with the service provider to ensure integrity and prevent misuse.

Consent History (Agent): Data is stored locally on the user's device, already reducing the risk of misuse. However, misuse may occur through the user themselves or malicious third party tools.

Unauthorized access must therefore be prevented by technical measures in line with the state of the art (see 6.2.2).

6.3 Consent History: Confidentiality and Availability

This section is non-normative

Discussion on the confidentiality and availability of the consent history (Agent):

In the current design, consent history data is stored unencrypted in the browser extension's local storage. This provides users with direct, local control and reflects the similar sensitivity and structure of consent history and browsing history.

The data is not encrypted in order to maintain usability. If encryption were implemented, users would need to enter a password on each visited website to enable communication with consent banners, which would likely increase consent fatigue rather than reduce it.

Because the data is stored unencrypted, it may be accessible to malicious browser extensions. This risk is considered limited, as it is comparable to that of unencrypted browsing history and is restricted to other installed extensions, which users and browser vendors can monitor and control through extension management and review processes.

Encrypting local data does not entirely mitigate risk of disclosure introduces further risks (weak passwords, loss of access if passwords or devices are lost, and exposure to brute-force attacks). Some of these risks may be better mitigated in certain cloud-based architectures, where repeated access attempts can be detected and blocked.

On this basis, the current solution stores consent history unencrypted in local browser storage, with the understanding that this can satisfy applicable accountability and security requirements. Further measures, such as encryption or alternative storage models, are considered potential future enhancements in line with evolving best practices.

6.4 Signal integrity

Signal integrity is intended to ensure that the consent actually given by the end user cannot be altered by the service provider or any third party.

The protection objective is implemented in an initial step through a signature procedure. For this purpose, the consent banner transmits the user's consent to the backend system, where the consent is assigned a unique Consent ID, digitally signed, and then returned as a Consent Record to the consent banner. The banner subsequently distributes the signed Consent Record across different channels (cookie, consent store, and consent history within the agent).

Only when such a signed consent has been received in the service provider's consent store are the subscribed third-party services, and thus data collection, activated. The digital signature ensures that the consent cannot be altered afterward without detection. This approach represents an effective implementation of Article 32 GDPR and reflects the current state of the art.

6.4.1 Cryptographic Signature

Non-repudiation

The current specification of consent records does not yet provide full technical guarantees of their immutability. The following section therefore introduces initial steps towards improving immutability by enabling the signing of consent records. However, this represents only one component of achieving fair non-repudiation—meaning that neither service providers nor users can credibly deny that a specific consent record was sent and received. In the current solution, this assurance is provided exclusively through organisational measures. To achieve fair non-repudiation on a technical level, the involvement of a trusted third party would be indispensable (see also ISO/IEC 13888-1:2020).

The server creates a cryptographic signature for each consent record using an elliptic curve digital signature algorithm (ECDSA) with SHA-384 hashing. The signing key is stored in a hardware security module that prevents the key material from being extracted.

The signature format includes a key identifier prefix that allows for key rotation. If a new signing key is deployed, the identifier changes, allowing the system to know which key was used to create each signature.

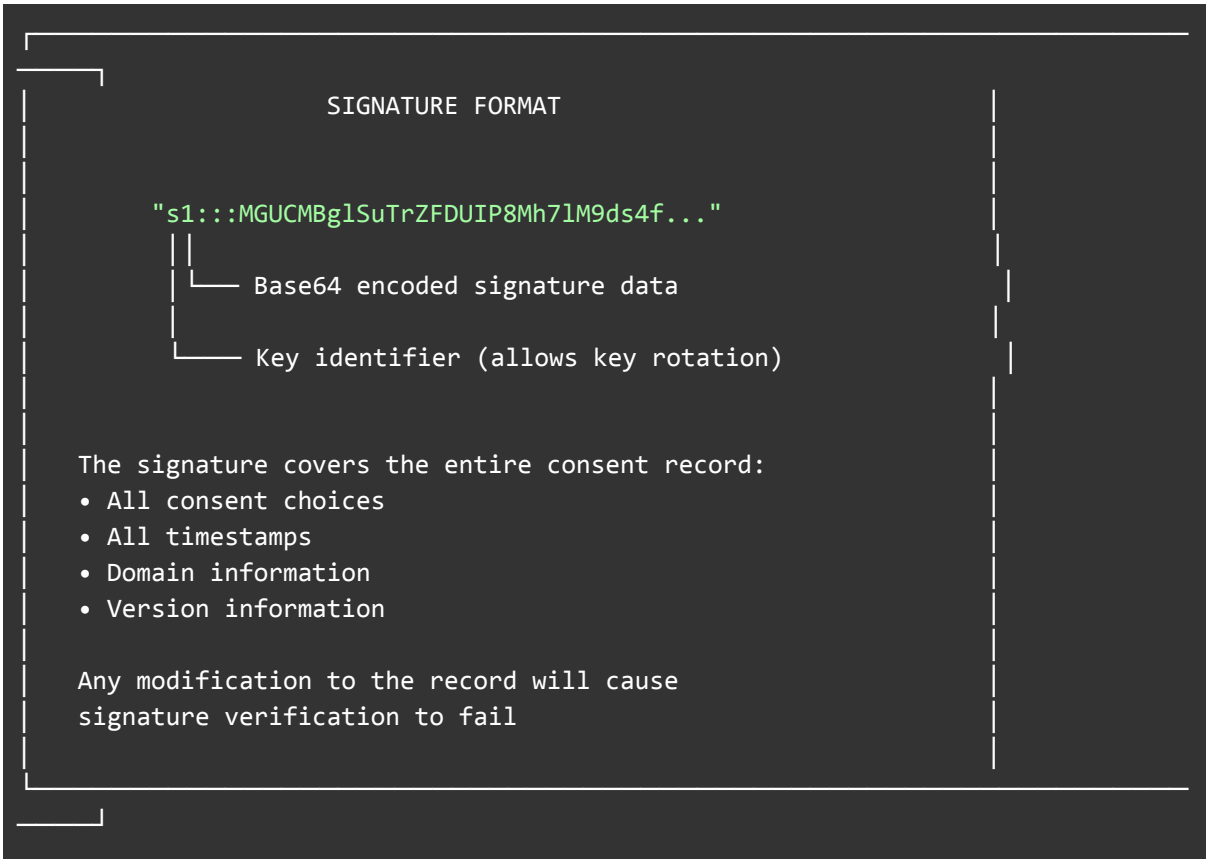

6.4.2 Limitations / progressing state of the art

The procedure has remaining limitations. If the service provider operates the consent store itself, i.e. signs and stores the consents itself, there is no direct control by a third party who, as a trusted third party, ensures compliance with the requirements of Articles 25 and 32 GDPR. However, if such a trusted third party is available on the market, this represents the state of the art, which the service provider must take into account.

However, even the trusted third party solution outlined here still has limitations. This is because it cannot yet prevent a service provider or an integrated third-party service from generating a false

consent, submitting it to the backend system of the trusted party for signing, and reintroducing it into the described consent channels. This risk can, indeed, be addressed through legal and organizational measures: service providers are required, by contractual obligations, to adhere to explicit prohibitions and due diligence requirements regarding integrated third-party services. In addition, a technical audit mechanism can be established. Because each signed consent is also stored in the user's local cookie, sample-based verification can be performed through an auditing frontend agent, which compares stored user settings against reference configurations. This agent can also observe network traffic to verify which data a service provider (or an integrated third-party service) transmits within the system and to whom. These measures represent a further enhancement of Article 32 GDPR compliance and an advancement of the state of the art.

With the spread of the EUDI Wallet, further advances in state-of-the-art technology are to be expected. To reduce central data storage, local data retention on the user's device presents a viable approach. The key challenge remains to ensure adequate protection of this locally stored data against unauthorized access or disclosure. Future development may therefore include local data storage within an encrypted database operated by a native application on the user's device. This architecture, currently gaining prominence through initiatives such as the EUDI Wallet, can be referred to as a "Consent Wallet." Within such an application, cryptographic keys—for example, for asymmetric signing of consent records by the user—could also be securely maintained.

Functionality of the state of the art requirements in Articles 25 and 32 GDPR
The above explanations are a good example of how the dynamic reference in Articles 25 and 32 GDPR to the current state of the art, which a service provider must take into account, works. As soon as a solution is available on the market that enables more effective protection against the risks of data processing for the data subject, service providers must take this into account. As a rule, this only does not have to be implemented if the costs are disproportionate. The service provider's resources are one important factor in this balancing exercise.[33]

7. Objection and withdrawal mechanism (Opt-Out)

Under Article 21(5) GDPR, controllers are required to honor opt-out signals submitted automatically by data subjects through technical specifications when processing is based on legitimate interests under Article 6(1)(f) GDPR. Based on the W3C's Global Privacy Control (GPC) this specification defines such a signal, allowing users to indicate that they do not wish their data to be processed for legitimate interests and to withdraw their previously given consents.

Sidenote: If GPC were implemented as the sole solution for automated signaling of user decisions, consent rates for controllers would likely decline sharply. This is because the signal currently provides only a technical means to automatically object to processing or withdraw previously given consent—without offering any mechanism to grant consent automatically. Consequently, a broadly supported specification for exchanging consent signals between stakeholders is essential, also from the controllers' perspective.

[33] EU Commentary Articles 25 and 32, Lee Bygrave, Paul de Hert…

7.1 Building on Global Privacy Control (GPC)

GPC defines a signal that allows users to express their preference not to have their data sold, shared with third parties, or used for cross-context behavioral advertising.[34] Under GPC the browser does two things:

(1) Adds an HTTP request header

```
Sec-GPC: 1
```

(2) Exposes a JavaScript property

```
navigator.globalPrivacyControl
```

The use of HTTP headers has the key advantage of timing: the server receives and processes the signal before any page content is rendered or scripts executed, ensuring enforcement of the user's pre-settings from the initial point of contact, which enables controllers to ensure that no non-essential cookies or trackers are initialised when a user visits the website.

GPC transmits binary signals 1 (opt-out) or 0 (no opt-out) for every website visit in order to object to the processing of personal data for the purposes predefined in the California Privacy Rights Act (CPRA) and California Consumer Privacy Act (CCPA).

If the header is absent entirely, no signal is considered set.

The header is transmitted on **every single request** — page loads, subpages, images, API calls — not just the first visit.

```
http
GET / HTTP/1.1
Host: example.com
Sec-GPC: 1
```

```
Two possible states:
```

```
Sec-GPC: 1   → user has opted out
Sec-GPC: 0   → no opt-out
```

7.2 HTTP Header Injection

This section describes how the consent agent (browser extension) injects a GPC-like binary signal, Sec-GPC: 1, into the HTTP headers of outgoing requests. Unlike the existing GPC specification, this signal is not intended to express an opt-out from the sale or sharing of personal data or from cross-context behavioural advertising, as these processing purposes require informed consent (opt-in) under the GDPR in any event. Instead, the signal indicates that the user wishes to

- opt out of any processing based on legitimate interests under Article 6(1)(f) GDPR, and
- withdraw all previously given consent.

---

[34] Human et al., https://www.dataprotectioncontrol.org/spec/, accessed 1st April 2026.

For HTTP header injection, the consent agent uses the browser's declarative network request API to modify outgoing HTTP requests via predefined header-modification rules.

Example:

```
{
    id: 10_000,
    priority: 1,
    action: {
        type: "modifyHeaders",
        requestHeaders: [{
            header: "Sec-GPC",
            operation: "set",
            value: "1",
        }, ],
    },
    condition: {
        resourceTypes: ["main_frame"],
    },
}
```

⚠️ **Limitations**

The HTTP header injection described here has so far only been tested in Firefox and Chrome. It is currently not supported in Safari in this form.

This signaling is currently not fully integrated into the ConStand ecosystem—including the Consent Banner, Consent Agent, and Consent Record Store—because of limited browser support. The navigator.globalPrivacyControl API is currently only implemented in Firefox, Brave, and DuckDuckGo, and is not supported by Chrome, Safari, Edge, or Opera. How full integration into ConStand would look is outlined in section 5.3 below.

7.3 Full integration into ConStand (Outline)

This section describes how, in theory and assuming full browser support, GPC-based opt-out signaling would be fully integrated into the ConStand ecosystem, including all components such as the Consent Banner, Consent Record, and Consent Store.

Even though opt-out mechanisms do not have to meet the same stringent standards as informed consent, the fundamental issues remain largely the same. Data subjects should be able to inform themselves about a processing purpose—whether framed as an opt-in or as an opt-out—before the processing takes place, that is, at a point when they have sufficient time to understand the associated risks and benefits. Similarly, objections communicated as opt-out signals must be properly documented so that both data subjects and controllers can verify whether an objection was (or was not) present at a specific point in time.

For this reason, achieving full integration of opt-out signaling into ConStand should be a key objective—not least because Article 21(5) GDPR requires controllers to enable objections through automated means. Realizing the GPC-based consent flow described in this specification, however, presupposes active involvement of key stakeholders such as browser vendors and an adaptation of GPC itself to EU-specific legal and technical requirements.

The website (or CMP on behalf of the website) can use the `navigator.globalPrivacyControl` API to receive GPC opt-out signals and store them in the consent store, as specified in 3.2 and 4.4.4.

⚠️ **Caution:** The navigator.globalPrivacyControl API is currently implemented only by Firefox and Firefox for mobile. To facilitate the incorporation of opt-out signals into the process proposed in this specification, all browsers should be required to implement this API so that the signal can be reliably received. This could call for corresponding legislative measures as part of the proposed digital omnibus under Article 88b.

Including GPC into the ConStand signaling could work as follows:

7.3.1. Detecting GPC signal on page load

On every page load (or at the earliest point your CMP script runs), read the browser API:

```
js
const gpcEnabled = Boolean(
  navigator.globalPrivacyControl === true
);
```

- `navigator.globalPrivacyControl` reflects the `Sec-GPC: 1` header:
- `true` → user has sent a GPC opt-out signal.
- `false` or `undefined` → no explicit GPC signal.

7.3.2 Mapping of the GPC signal to the consent model of the CMP

⚠️ **Caution:** GPC in its current form is limited to a binary “do not sell/share” signal, which does not reflect purpose-specific objections as required by Article 21 GDPR. In this regard, the GPC signal must be extended so that it can transmit purpose-specific signals according to the respective purpose taxonomy referred to section 2.4.

Treat the GPC signal as a **pre-existing objection** (e.g., “do not sell/share”) for relevant categories as specified by the respective CMP (e.g. advertising/analytics/vendors).

If GPC is `true` and no existing consent exists for the user, set those categories to rejected by default.

7.3.3 Storing of the signal in the consent store

Store consent record in the consent store, cookie and consent agent in accordance with 3.2 and 4.4.4. The record should include:

- The GPC-derived state (e.g. `ads: "rejected"`).
- A metadata field indicating that the decision came from `navigator.globalPrivacyControl` (for logging and auditability).

Example structure (simplified):

```
json
{
  "timestamp": "2026-03-26T15:00:00Z",
```

```
  "gpc_detected": true,
  "gpc_original_value": true,
  "preferences": {
    "necessary": "allow",
    "analytics": "reject",
    "advertising": "reject"
  }
}
```

7.3.4 TPP blocking

The service must block the collection of personal data by any third party service provider immediately or prevent third party scripts from loading in the first place - in accordance with the signals.

Blocking of third-party services can be handled in accordance with the regular consent process as defined in section 4.3.2.

7.3.5 Re-consent

If users want to consent to a purpose for which they have transmitted an opt-out signal before, this should be handled as a regular opt-in signal in accordance with the process outlined in section 4.3.2.

8. Version management

8.1 Why Version Management Matters

❓See for an in depth description of the versioning and its underlying risk assessment: Annex 3.

The controller's processing operations and consent banner configuration usually change over time. New purposes may be added, existing purposes may be modified, or services may be added or removed. When these changes occur, it may be necessary to notify users or ask them to reconsent.

Version management ensures that users are appropriately informed about changes while avoiding unnecessary interruptions for minor updates that do not affect their choices.

When a controller alters its data processing operations—such as by changing data categories, storage locations, personalization models, or third-party services—this can increase risks to data subjects' fundamental rights. Since data subjects base consent decisions on these risks (see 5.2.2.1), later data processing operations which lead to higher or even other risks cannot anymore be based on the original consent. This means that, according to the case, the controller must obtain consent anew or at least notify users about the changes, offering them a chance to withdraw.[35]

Thus, if a controller makes significant changes to processing operations after a data subject has consented, the original consent may no longer serve as a valid legal basis. The system must

[35] This results from the purpose limitation principle according to Art. 5(1)(b), 6(4) GDPR.

therefore detect differences in risk levels between the time of consent and the data subject's subsequent return to the service.

Example 9
User U visits publisher P's website in January and consents to personalised advertising to access P's content. U gives this consent on the basis of a comparatively favourable risk profile, as P's system collects behavioural data, aggregates them into large advertising profiles, and shares them with an ad network of around 30 publishers within the EU.

In March, P seeks to increase ad revenue and therefore joins a significantly larger ad network comprising more than 900 companies inside and outside the EU, all receiving personal data collected from P's users. Because many more companies now have access to U's data, the risks to U's privacy and autonomy increase.

When U returns to P's website in April, the underlying risk profile has changed. This requires an assessment of whether these changes remain covered by U's original consent, i.e. whether the new processing is still compatible with the purpose for which U initially consented (purpose compatibility assessment). This assessment is carried out by comparing the triggers of both states, and if a higher-risk trigger is identified, the data subject is notified and given the opportunity to opt out of this higher-risk processing.

The following version management system allows controllers to indicate relevant changes to their processing operations and their associated fundamental rights risks and compare them to the state where each individual consent record was created i.e. where consent was given or refused by an individual data subject. This creates an automated purpose compatibility assessment taking place each time a user returns to the specific service of the controller.

The system includes several triggers that the controller may activate when modifying processing activities. Each activation is logged in a trigger count, which forms the basis for initiating corresponding actions for returning users.

It is the controller's legal responsibility to activate these triggers accurately and in accordance with the actual changes made to their processing operations. The provider of the consent banner may automate this trigger mechanism to prevent incorrect or misleading activations by the controller. However, such automation falls outside the scope of this specification, which focuses solely on the communication of signals between multiple stakeholders.

Examples:

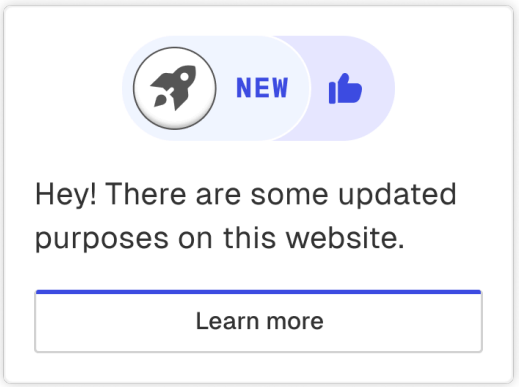

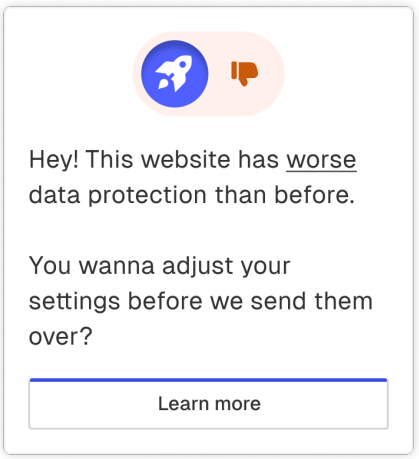

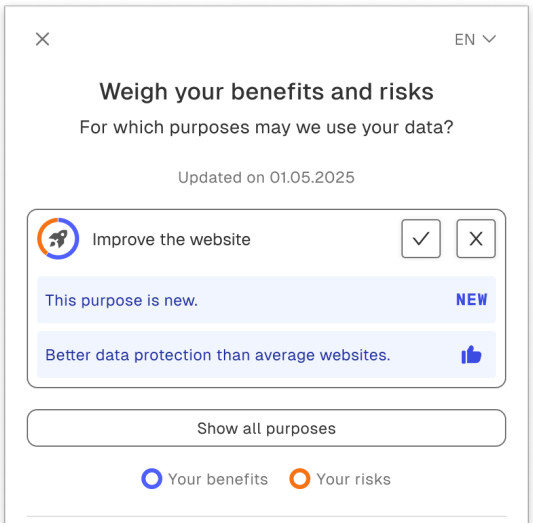

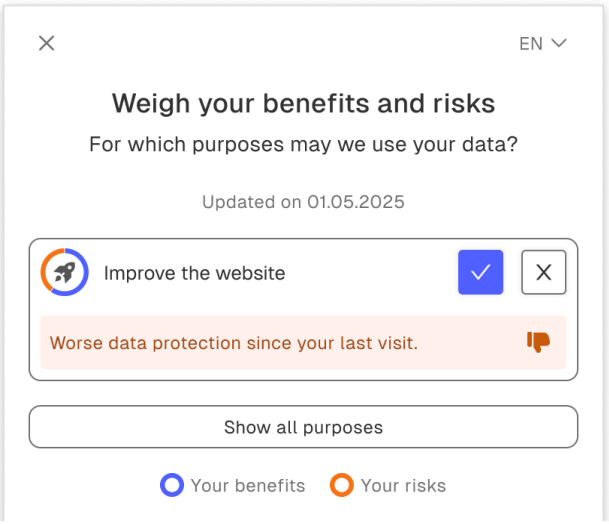

8.2 The Trigger System

Each purpose has an associated trigger number in the format "A.B.C" where each component represents a different type of change:

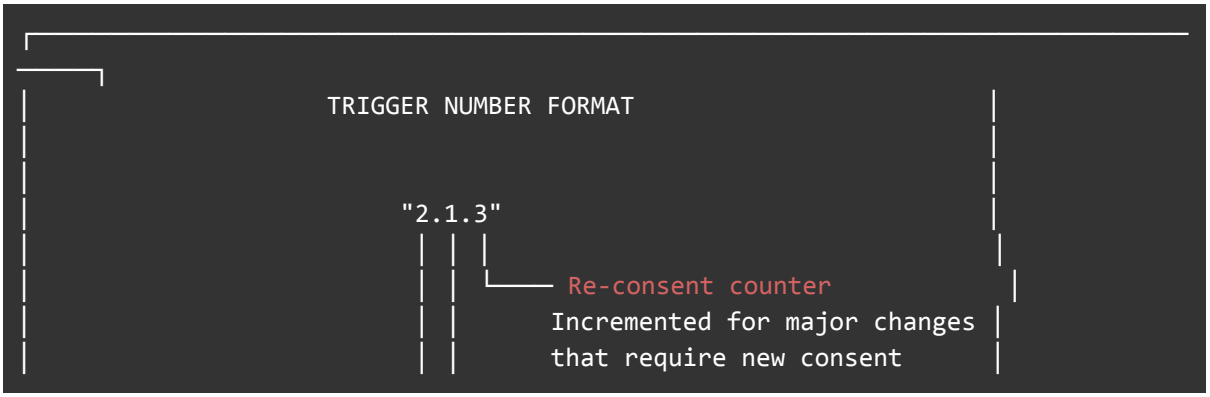

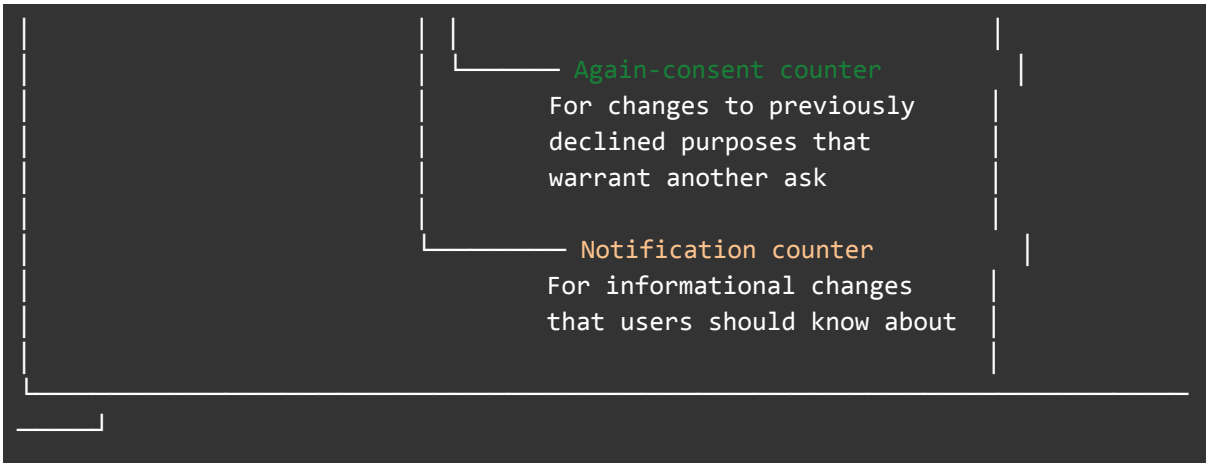

8.3 Response to Version Changes

Based on the comparison, the system takes one of several actions:

No action needed when the trigger numbers have not changed or the changes do not affect the user's consent state. The user's existing consent is applied without interruption.

Notification: Show notification when changes occur that users should be informed about but that do not invalidate their existing consent. The notification briefly explains what has changed and provides an option to review the details and, if desired, withdraw the still-valid consent.

Again-consent: Prompt the user to reconsider consent when a purpose they previously declined now presents a significantly lower risk than before. The user is shown the updated purpose information and given the opportunity to provide consent again.

Re-consent: Require re-consent when major changes mean the user's previous consent no longer applies. The full consent interface is shown and the user must make new choices.

8.4 How Version Comparison Works

When a user returns to a website, the system compares the trigger numbers in their consent record against the current configuration.

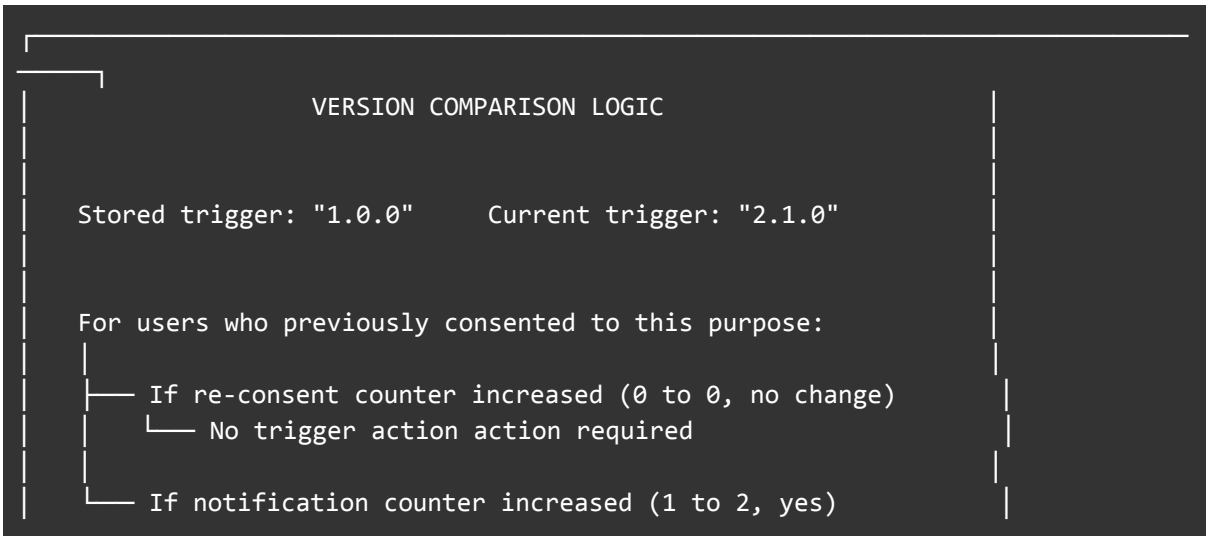

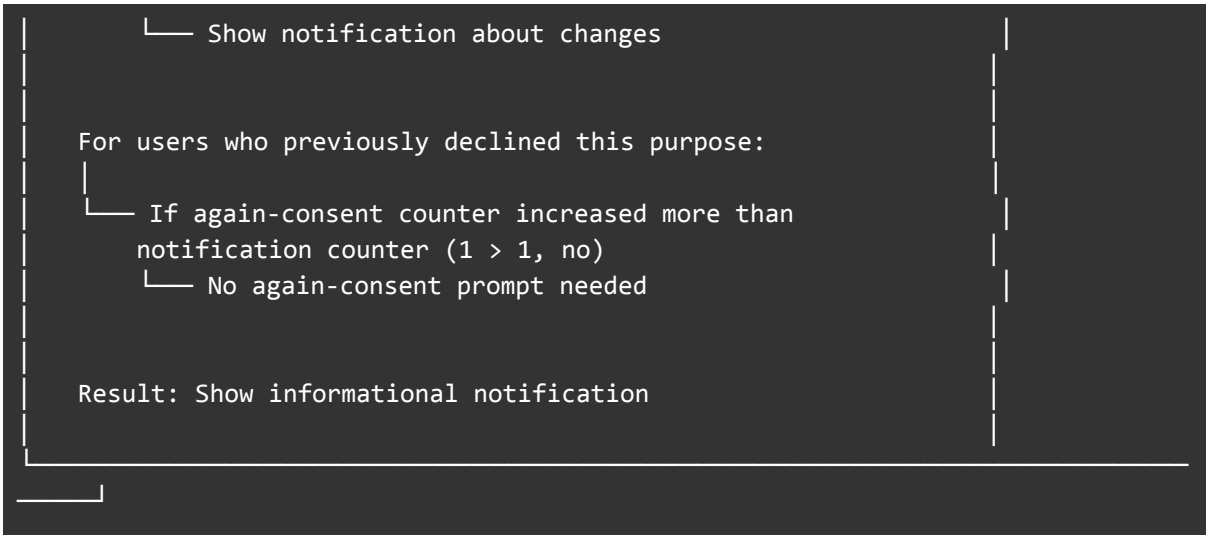

9. Extension to Mobile Apps and other contexts

Summary
In the context of mobile apps, the focus of the consent process lies in transmitting users' consent decisions to the consent agent, ensuring they have a centralized access point to review and modify their choices for a specific app in the future.

As personal data is collected in similar ways within mobile apps as it is on websites, data processing in mobile environments equally requires informed user consent. While integrating mobile apps into the consent process, as described in this specification, introduces certain challenges due to the closed and proprietary nature of app ecosystems, such integration remains feasible through several approaches.

A key difference in the mobile context concerns the need for automated transfer of consent signals from the consent agent to the app (e.g., its consent banner). Mobile apps do not necessarily require such automated transfer of pre-settings, since consent is typically obtained at the moment users download or first open the app. As a result, consent fatigue — a major concern in web environments — is generally less relevant for mobile apps.

The primary focus of integrating mobile apps into the consent process is to ensure that users have a centralized overview of their past consent decisions, enabling them to review or modify these choices at any time — for instance, to withdraw previously granted consent or to grant consent that was previously refused. To achieve this, consent decisions made within a mobile app must be transmitted to the consent agent, which maintains the central consent history.

There are two main ways in which the consent process can be implemented:

9.1 Integrating the consent design and API directly into the app

This approach means the consent interface and logic (how consent is requested, stored, and communicated) are embedded within the mobile app itself.

- How it works:
  The consent authority (CA) provides an SDK (Software Development Kit) that app developers integrate into their app. The SDK includes the visual components (e.g. consent pop-ups, toggles) and the necessary APIs to send consent decisions securely to the CA.

- Advantages:
    - Seamless and fast user experience within the app (no need to leave the app).
    - Can be customized to match the app's design and branding.
- Challenges:
    - Requires technical integration efforts by each app developer.
    - May reduce the perception of independence, since the consent interface appears as part of the app rather than a neutral third party.

Example 10
A fitness app includes a pre-built consent screen from the consent agent's SDK. When the user accepts or refuses, the app sends that data directly to the consent store.

9.2 Guiding users to the CA's website during installation

In this approach, the consent process happens outside the app, on a trusted external website managed by the Consent Agent provider.

- How it works:
  When users install or first open the app, they are redirected (for instance through a link or system prompt) to the consent agent's web portal. There, they review and consent to the app's data usage policies.
- Advantages:
    - Strengthens user trust: consent is collected by an independent third party, not the app publisher.
    - Reduces development workload for app creators since they only handle the redirection.
- Challenges:
    - Requires a short user flow interruption (users temporarily leave the app).
    - Needs coordination between the app and CA systems to pass back consent confirmation.

Example 11
When installing a public transport app, the user is directed to the consent agent's secure site to provide or refuse consent. The CA then sends a confirmation token back to the app to proceed with the data processing to which the user has provided consent.

9.3 Outlook: Further contexts

This approach could also be applied in other contexts where consent forms the basis for data processing. One such context is the doctor's office, where patients routinely sign privacy agreements manually. In this setting, informing patients about the use of their data in connection with medical treatment can be achieved more effectively through contextual consent agents than solely at the time of an in-person visit.

10. Treatment of non-conformity

If controllers do not comply with this specification, because they

- do not accept signalling in accordance with this standard,
- do not support the included risk assessment or version management,
- do not support the purpose base layer and do not provide a mapping of their custom purposes to this base layer,

consent banners appearing on the services of such controllers may be blocked in order to prevent consent fatigue from persisting.

Further advancement

If non-conforming controllers request consent for purposes that are not supported or not mapped in accordance with this specification, consent agents may, where feasible, perform the corresponding purpose mapping themselves.

Where such purposes do not yet exist in the system, controllers should be able to apply for the addition of new purposes so that emerging personal data processing practices can be adequately reflected.

Annex 1: Purposes and Mapping

1. TCF Bibliography

- **Purposes:** High-level reasons for processing (e.g. storing data, personalising ads). Vendors declare them, and users can give or refuse consent per purpose.
- **Special Purposes:** Essential processing activities (e.g. security, basic delivery) that do not require opt-in but must be disclosed to users.
- **Features:** Technical methods that support one or more purposes (e.g. linking devices). They do not have their own legal basis and inherit it from the purposes they serve.
- **Special Features:** Particularly sensitive technical capabilities (e.g. precise geolocation) that require explicit opt-in consent.

2. Purposes (current set)

1. Store and/or access information on a device
2. Select basic ads
3. Create a personalised ads profile
4. Select personalised ads
5. Create a personalised content profile
6. Select personalised content
7. Measure ad performance
8. Measure content performance
9. Apply market research to generate audience insights
10. Develop and improve products
11. Use limited data to select content

3. Special Purposes

1. Ensure security, prevent fraud, and debug
2. Technically deliver ads or content
3. Save and communicate privacy choices

4. Features

1. Match and combine offline data sources
2. Link different devices
3. Receive and use automatically sent device characteristics for identification

5. Special Features

1. Use precise geolocation data
2. Actively scan device characteristics for identification

6. Exemplary purpose mapping

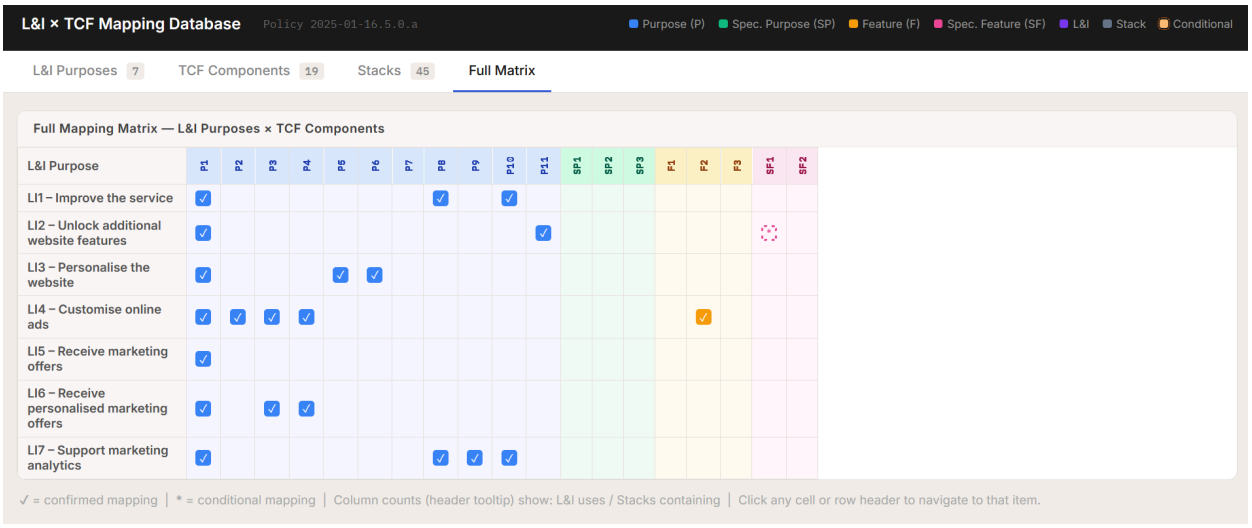

Annex 2: Navigator.consent API for communication between agents and banners

⚠️**Challenges and Limitations**

Navigator.consent remains a draft specification requiring native browser implementation to function as a standardized communication layer between CMPs and consent agents. This process involves advancing through standards bodies such as W3C, WHATWG, or WICG.

Critical limitation: The current API lacks support for essential data needed for a complete consent procedure as outlined in this specification—particularly the communication of risks and benefits (see 5.2.2.1) and versioning parameters (see 8.).

Ensuring interoperability between consent banners on the controllers' side and consent agents on the users' side is essential for achieving seamless automation and consistent management of consent decisions across all stakeholders and throughout the value chain.

For CMPs — and consequently the consent banners used by their customers — to properly respect consent signals from external consent agents, all parties must be able to communicate seamlessly.

Today's consent banners lack a standardized or direct interface for such communication. As a result, consent agents are forced either to block these banners entirely or to scrape each website's DOM to extract consent-related information. This, in turn, leads to complex and often inaccurate purpose-matching procedures that attempt to automatically "fill out" the various CMP banners on behalf of users.

In order to solve this communication issue, this specification could connect to and build up on the navigator-consent API,[36] which proposes a browser API as transport and coordination layer between CMPs and Consent Agents. Navigator.consent would replace the windowAPI currently used in this spec.

- The proposed API carries structured preferences per vendor and per purpose.
- CMPs can both read existing preferences and receive updates from assistants, and assistants can listen for consent events and send updated preferences back, all via the same, defined interface.
- The CMP still owns storage, scope (per site, per domain, etc.), and compliance records; the browser API is only a transport/coordination layer.

The inclusion of such a Browser API as standardised communication layer between CMPs and consent agents would work as follows:

1. CMP Integration

CMPs declare themselves in the DOM context after loading via:

```
navigator.consent.registerInterface({vendor: "your-cmp-name", versionIdentifier:
"1.0", regulation: "gdpr", jurisdiction: "DE"}).
```

This call returns a registrationId required for subsequent operations.

2. Vendor and Purpose Registration

[36] See https://www.navigatorconsent.org/rfc.

CMPs register vendors using

```
registerVendors([{id: "google-analytics", name: "Google Analytics", domain:
"google-analytics.com", privacyPolicyUrl: "https://..."}]).
```

CMPs register purposes with

```
registerPurposes([{id: "analytics", name: "Analytics", legalBasis: "consent"}]).
```

3. Preference Synchronization
CMPs synchronize existing preferences by calling

```
updatePreferences({vendors: {"google-analytics": "grant"}, purposes:
{"analytics": "grant"}, source: "cmp"}).
```

4. Consent Request Flow
To notify consent assistants and apply user preferences, before displaying the consent banner, CMPs call:

```
requestConsent({registrationId, vendorIds: ["new-vendor"]})
```

5. Event Handling
CMPs listen for updates from assistants via:

```
navigator.consent.addEventListener("update", (event) => { /* handle changes */
}).
```

6. Data Transmission
CMPs provide the following data through registration and update methods:

```
Metadata: CMP details (ID, version, regulation), vendor lists (ID, name, domain,
policy URL), purpose lists (ID, name, legal basis).

Preferences: States ("grant", "deny", "unset") for vendors/purposes, with
reasons and timestamps.

Requests: Scoped consent needs or withdrawals.
```

All payloads must be JSON-serializable objects.

7. Further information

CMPs Cannot Read Assistant Data or Modify Other Registrations
Consent Management Platforms can only write their own data (vendor lists, preferences). They cannot:

- Access preferences set by consent assistants (browser extensions).
- Edit or delete registrations from other CMPs/assistants.

This prevents CMPs from spying on competitors' consent states or interfering with user choices made elsewhere.

Browser Audits Enforce Compliance
The browser logs all API calls with metadata (who called what, when, context). This audit trail ensures accountability—violations can be detected during compliance reviews.

Context Boundaries & NotAllowedError
The API runs in two isolated contexts:

- DOM context (CMP script on the webpage).
- Extension context (privacy assistants).

Extension-only methods (like reading stored preferences) cannot be called from webpage JavaScript. Attempting this throws NotAllowedError, preventing websites from bypassing extension isolation.

Example: The CMP calls requestConsent() → assistant applies its prefs → CMP reads the update event (allowed). But getPreferences() from DOM → error (blocked).

Annex 3: Automating risk assessments on websites

The (third-party) **tools used** on websites and **how to configure them** directly impact data protection risks for website visitors. If high standards of data protection are set, this can be made visible to website visitors.

1. Risk score: How does it work?

When a controller configures a consent banner in the Banner Configuration Panel, the selections feed directly into an automated risk assessment. This assessment calculates a risk score based on the configurations in the customer panel, reflecting the potential impact of the data processing on website visitors' fundamental rights. It evaluates factors such as data processing purposes (of the controller and third parties), data categories collected, storage location and duration, tracking methods used, and—where applicable—personalization models employed.

Each factor is assigned a weighted value reflecting its relative impact on the rights and freedoms of data subjects. The resulting weighted calculation produces a specific score for each risk arising from the processing of personal data, ranging from 0 (no risk) to 100 (highest risk) per fundamental rights risk. To make this accessible for laypeople, this score is translated on the visual end-user interface into a **three-dot system** (low, medium, high), displayed clearly to website visitors. Besides the risks, also the benefits are communicated to website visitors.

For each risk, the resulting score is benchmarked against a typical baseline for average websites. This makes transparent whether the controller's practices pose higher or lower risks to the fundamental rights of users compared to an average website.

The score helps controllers to understand the privacy impact of their technology choices and can motivate them to adopt more privacy-friendly tools or configurations. The users, in turn, gain clear insight into the controller's efforts to safeguard their personal data. When improving the setup and thus decreasing fundamental rights risks, website visitors are discreetly informed, strengthening trust and, over time, increasing consent rates on the website.

2. Trigger System: How does it work?

The so-called "trigger system" integrates with the risk calculation mechanism (risk score) by adding an enforcement layer.

- It ensures that users remain informed about relevant changes to the technical system.
- It monitors that data processing is only carried out with effective consent.
- It enables website providers which improve data protection for their users, to ask for consent again, thereby boosting consent rates and creating a competitive advantage for sites with high privacy standards.

2.1 Making changes to the processing operation transparent to users

In the context of consent-based processing of personal data, changes to the technical system may change the risk profile, and thus the knowledge base on which website visitors relied when deciding whether to give or withhold consent. This can have an impact on the validity of prior consent decisions and must be managed accordingly. When controllers makes privacy-enhancing improvements to their system, informing users about these changes can help strengthen their trust in the organization. However, when introducing new technologies that are more privacy-intrusive, controllers are required to inform users accordingly.

Example:
Consent can be thought of as a contract between website visitors and controllers, granting controllers the right to process personal data of website visitors.

In any contract, both parties must understand what they are agreeing to, including the fundamental terms. If one party unilaterally changes those terms, they must—at minimum—inform the other party, or, depending on the severity, formally renew the contract.

Data protection consent follows the same principle. If users consented based on a low-risk profile, they must be informed (at least) if risks increase significantly. Consenter enables this for the first time—setting a new state of the art in consent management.

As Article 25 GDPR requires controllers to implement state-of-the-art measures, website providers must manage consent accordingly for full GDPR compliance.

This creates the following scenarios:

Again-consent: If the risks decrease significantly, the basis for previous refusals may have changed. In such cases, controllers may request consent again from users who had previously declined, likely boosting overall consent rates.

Notification: If risks increase significantly, controllers must inform the users who had consented before on the basis of a lower risk profile, giving them the chance to opt-out of the riskier processing.

Re-consent: If the changes introduce entirely new risks, this is typically tied to a purpose change. Prior consent of website visitors in this case does not serve as valid legal basis for the processing of personal data. For this new purpose controllers must obtain fresh consent.

Example:
When using Google Analytics solely to improve the website based on usage data, consent must be obtained specifically for the purpose of “website improvement.”

If data transmissions to Google for their own advertising purposes are activated, this introduces new risks for website visitors. In that case, controllers must obtain fresh consent for the corresponding advertising purposes.

2.2 Automation

The so-called “trigger system” (see 8.2) automates these actions (again-consent, notifications and re-consent). The risk calculation mechanism (see above) detects risk increases or decreases as controllers adjust settings in the Banner Configuration Panel. Before saving, controllers are prompted

to **trigger** re-consent, notification, or again-consent as needed—or do nothing. The controllers are responsible for entering all information truthfully, including trigger selections.

If users who previously consented return to the website, the trigger system compares the triggers from the time of their consent against the triggers of the current configuration. If the risks—and thus the triggers—have changed in favor of or to the detriment of the data subject, users are informed accordingly or asked for consent again. This is achieved by versioning each consent decision and cookie banner configuration.

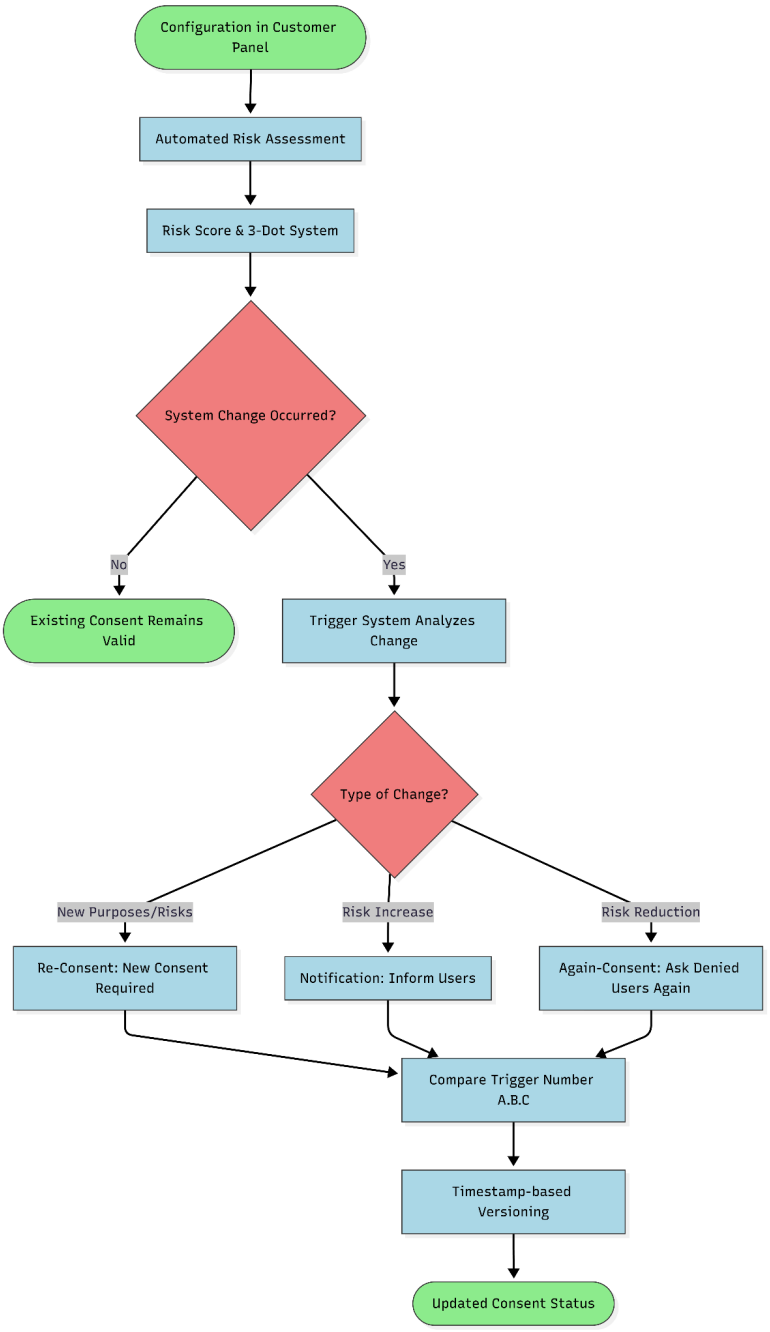

3. Risk assessment in detail

3.1 Risks and baseline floors

Each processing purpose generates distinct risks to the fundamental rights of affected data subjects. Since the purpose itself serves as the strongest indicator of these risks' severity, they are assigned a baseline value for each risk. No matter how privacy-friendly controllers configure the processing operation and their tools, the overall risk score for that purpose can never fall below this minimum threshold. However, if a specific risk fully ceases to exist—for example, because controllers no longer transmit personal data outside the EU/EEA—that risk disappears entirely.

Purpose	Risk	Baseline
Improve the service	Data transfer outside the EU	0
	Statistical insight into how you use the website	0
Unlock additional website features	Data transfer outside the EU	0
	Insight into the use of the additional function	20
Personalise the website	Data transfer outside the EU	32
	Profiling of your interests	20
	Influencing how you use the website	50
	Higher prices	32
	Different display of the website	32
Customise online ads	Data transfer outside the EU	32
	Profiles of your purchasing behavior also by third parties	50
	Motivation to buy products	50
	Higher prices	65
	Restricted product range	50
Receive marketing offers	Data transfer outside the EU	0
	Motivation to buy products	20
	Direct approach	0
Receive personalised marketing offers	Data transfer outside the EU	33
	Profiles of your purchasing behaviour	20
	Motivation to buy products	50
	Direct approach	20
	Restricted product range	0
Support marketing analytics	Data transfer outside the EU	0
	Statistical insight into how you use the website	32

3.2 Weighted risk values

Based on the floor values introduced before, the following choices in the Banner Configuration Panel play into the calculation of the risk score:

3.2.1 Tracking method

One tracking method must be selected for each **data recipient**.

Tracking method	**Value**
No tracking	0
First party tracking / single session	2
Cookieless First party tracking / cross session	2
First party / Cross session tracking	6
Third Party (Cross website/single session) tracking	8
Third Party (Cross website/cross Session) tracking	12
Third Party (Cross website/cross Session/cross device) tracking	14

3.2.2 Legal role

One legal role must be selected for each **data recipient**.

Legal role	**Value**
Self-hosted	2
Processor	4
Joint Controller	6
Controller	10

3.2.3 Personalisation

One personalisation model must be selected for each **data recipient**.

Personalisation model	**Value**
No personalisation	0
Group based (properties)	2
Group based (behaviour)	4
Profile based	10

3.2.4 Data category

One or more storage locations must be selected for each **data recipient**. The total score is averaged across all data categories.

Data category	**Value**
Aggregated site statistics	4
Non-precise location data	4
Device characteristics	4
IP-Adress anonymised	4
IP address	8
Privacy choices	8

Authentication-derived identifiers	8
Browsing and interaction data	8
Device identifiers	8
User-provided data	12
Precise location data	12
Probabilistic identifiers	12
E-commerce Activity	12
Direct identifiers	12
Social media interaction data	12
Special categories Art. 9 GDPR	20
Users' profiles	20

3.2.5 Storage duration

One storage duration must be selected for each **data category**.

Storage duration	**Value**
< 1 year	-1 (per data category)
> 1 year	1 (per data category)

3.2.6 Storage location

One or more storage locations must be selected for each **data category**.

ℹ️ Data transfers outside the EU/EEA are treated as a distinct new risk, minimally impacting other risk scores.

Storage location	**Value**
EU	0
Adequacy decision	1
Appropriate Safeguards	2
Risk reduction per additional safeguard (floor is at 1)	1
No legal basis	50

3.3 System Overview