RETAIL ENERGY CODE

This document is classified as **WHITE**.
Information is for public, unrestricted dissemination and publication.

REC CONSUMER CONSENT TOPIC SPECIFIC SOLUTION GROUP
MEETING MINUTES

MEETING	**REC_CCTSSG_26_0604**
DATE	**04 June 2026, 10:00**
LOCATION	**By teleconference**

1. INTRODUCTIONS

The Chair welcomed attendees to the meeting of the Consumer Consent Topic Specific Solution Group (CCTSSG). This session would focus specifically on the Consent Record Schema and Security Profile.

The Chair advised that a poll will be issued to Members during the meeting to consider a new time slot for all future meetings going forward, noting significant clashes with other industry meetings at the moment. Members indicated their preference to move meetings to Wednesday afternoons, and RECCo agreed that this would take effect ahead of the next meeting.

The Chair noted that no pre-meeting papers were issued to Members ahead of the meeting, advising that due to the fast-paced nature of the project, some information and developments are not ready far in advance of meetings. The Chair added that RECCo will continue to provide Members with information on the topics for each meeting ahead of time, to ensure organisations can send the individuals with the most relevant expertise to each meeting.

2. CONSENT RECORD SCHEMA

GR explained that the Consent Record schema is a foundational component, defining how Consumer Consent is captured, structured, interpreted, and shared across the ecosystem, and therefore must be designed correctly. Key priorities include accurately reflecting consumer intent, ensuring strong data protection controls aligned to GDPR and related requirements (including transparency, data minimisation and informed consent), and maintaining sufficient flexibility and granularity to support future use cases.

GR clarified that the session should focus specifically on the schema itself (e.g. JSON/YAML structures within requests), rather than wider API specifications, although these are acknowledged as interdependent and will be shared subsequently. They noted that an initial draft has been developed with the technical partner but requires industry feedback, particularly from organisations already delivering similar services.

GR introduced the proposed use of Rich Authorisation Requests (RAR) as the consent model, explaining that this approach offers greater granularity than traditional scopes by capturing key elements such as what consent is for, who it applies to, its duration, and its purpose.

GR outlined a view of the FAPI 2.0 flow to demonstrate where the consent record and its associated schema are used throughout the process, noting that many technical components are linked to the consent object. They explained that the process begins with an Authorised Third Party (ATP) submitting a push authorisation request containing an initial RAR, signalling its intent to access data. CCS then generates a unique request identifier (request URI) and redirects the consumer to CCS, where they authenticate (and complete ID&V if required). At this stage, CCS enriches the original request with Consumer data (e.g. confirming property and MPxN association), forming a complete consent object. This is then presented to the Consumer for informed consent. Once consent is provided, CCS creates a consent record, issues an authorisation code, and subsequently a token, which the ATP uses to request data from Energy Data Holders (EDHs), who validate the token via introspection before sharing data.

Members raised concerns that the material was overly technical and lacked sufficient prior context, particularly regarding the Consumer journey and business requirements. One member challenged the proposed flow, stating that it appears inconsistent with real-world consent journeys. They highlighted that Consumers typically engage with ATPs through a clear value proposition, and that ATPs design journeys to help consumers understand why consent is required and to provide the necessary data in a user-friendly way. The member expressed concern that the model presented (where CCS derives and enriches data such as MPxN) felt disconnected from how consent is practically obtained and experienced, particularly given that many consumers would not know or provide such information themselves. It was noted that designing the schema in isolation from the user experience could be challenging. GR and SJ clarified that the session was intentionally focused on defining the content of the consent record itself—i.e. what information must be captured and presented to enable informed consent—rather than the wider journey. Participants were encouraged to review the material post-session and provide further feedback to Consumerconsent@retailenergycode.co.uk

GR presented the first draft of the consent record schema to Members. They emphasised that the schema presented is an early draft and under active development, forming the backend structure that both determines consent and reflects it to Consumers. The schema is based on RAR JSON and includes key fields such as data type and purpose, with ongoing consideration around the appropriate level of granularity and use of controlled (enumerated) values. They noted that an “other” category may be required for flexibility, with additional mandatory detail where used.

A discussion followed on how to handle multiple purposes within a consent. A member asked whether a single consent could cover multiple data types and purposes simultaneously. GR confirmed this is under consideration, including whether to structure this as multiple purposes within one consent record or as separate consent records. SJ highlighted that this decision is important for how consents are tracked and revoked, noting that separate purposes may need to be independently revocable even if presented as a single user journey. The member confirmed that in current practice, consent is typically withdrawn at the same level it was granted, though this may not be optimal. MPW added that consent structures may depend on whether a service requires multiple data inputs to function, while also noting potential complexities across fuels and use cases. The group agreed this is a design question to be

taken away, with options potentially allowing ATPs flexibility in how they structure and link consents.

A member raised concerns that including an “other” purpose field could undermine GDPR compliance, reduce clarity for consumers, and create challenges for machine readability, particularly in an automated, near real-time system. They suggested that, rather than allowing free text, purposes should be iteratively added through a controlled process. GR responded that while an “other” field could provide flexibility and avoid blocking ATPs, it introduces design complexity (e.g. how to ensure sufficient structure for data type, duration, and frequency). SJ explained that the inclusion of “other” was intended as a pragmatic fallback to avoid delays caused by formal change processes for adding new purposes. They noted that, if implemented, it would include mandatory additional detail for consumer transparency and would be monitored through assurance processes to ensure compliance and appropriate use.

A member raised the question of whether purposes should be standardised across all ATPs or allow for more bespoke, organisation-specific definitions. GR agreed that the solution should support both common and tailored arrangements, potentially through configuration or “domain authority” structures to reflect different data sharing models.

Another member recommended avoiding an “other” field altogether due to legal and governance risks, advocating instead for a robust but efficient change process with legal oversight. They also advised that consent should be stored as separate records where purposes are distinct, to support granular consent and revocation, even if presented to consumers in a bundled format. Additionally, they highlighted the importance of capturing the exact wording shown to consumers at the point of consent for audit and dispute resolution purposes.

A member emphasised the need to capture frequency as a distinct attribute within the schema, noting that consent use cases vary significantly—from one-off historical data requests to ongoing or near real-time data sharing. GR confirmed this is under consideration, with support for including frequency as a dedicated field alongside purpose.

GR advised that the schema design remains highly interdependent, with open questions across fields such as type, purpose, and frequency, and confirmed that current thinking broadly supports using data type and purpose as core elements to reflect existing and future data sharing arrangements. However, they emphasised that the schema must remain lightweight to deliver system efficiency, which is why certain elements (e.g. data controller or legal basis) are not currently included directly, even if captured elsewhere. A member responded that consumers should be able to see who the data controller is, particularly for complaints handling. GR agreed this is important but suggested it may sit outside the schema itself, provided it is accessible to consumers through other means.

GR sought views on whether to include default or open-ended end dates, and multiple members advised that consents should always have an explicit end date (to support GDPR compliance), but without a fixed default, as appropriate durations vary by purpose. Instead, responsibility should sit with the ATP to define appropriate durations, typically aligned to contractual or service needs (e.g. 12–36 months), rather than hard-coded limits.

One member highlighted the importance of future-proofing the design for cross-sector interoperability, drawing on experience from open banking where fixed consent expiry rules (e.g. 90 days) exist. They noted that different sectors have varying regulatory requirements and cadences for consent renewal, and that CCS must remain flexible to support cross-sector data sharing and differing validation requirements, while still aligning with GDPR. MPW acknowledged the need to balance energy sector requirements with broader cross-sector considerations and confirmed this will be taken into account in future design.

Further discussion explored how consent duration should be modelled. Members broadly agreed that using explicit “from” and “to” date fields is sufficient to represent both one-off and enduring consent, avoiding the need for additional fields. It was suggested that consent should always have a defined end date (rather than being open-ended), with ATPs responsible for determining appropriate durations.

GR asked Members how MPxN identifiers and fuel type should be represented within the consent schema. They noted that while Consumer ID&V is expected to confirm occupancy via address, it may not explicitly indicate whether gas or electricity data is being requested. As a result, RECCo are considering whether fuel type should be captured as a separate data field or instead be handled through greater granularity within the existing data type (i.e. within the data catalogue). A member responded that, in practice, the approach depends on scope. In smart metering contexts, MPXN identifiers already implicitly link to fuel type, whereas in broader scenarios this linkage may be less direct. They suggested keeping the model simple by focusing on clear identifiers and allowing the purpose and data structure to drive the required clarity, rather than adding unnecessary complexity.

The discussion then moved to purpose granularity. GR highlighted that more granular purpose definitions could help EDHs better understand exactly what data to provide and reduce the risk of over-sharing (supporting GDPR data minimisation). However, this introduces challenges given the absence of standardised data sets across the energy sector.

Participants broadly supported increased granularity. One member suggested separating concepts such as data type, frequency, and granularity to improve clarity, and another reinforced that highly specific purposes are important to ensure that data is only used for the consented activity, citing scenarios where organisations may initially request data for one purpose (e.g. bidding) but later require it for a different use (e.g. ongoing contract management). More granular purposes would help prevent inappropriate continued use of data beyond the original consent.

GR clarified that CCS will not define standardised data sets. Instead, EDHs retain control over their own data schemas, with the consent schema (particularly type and purpose) informing what data they interpret should be shared. The key challenge is ensuring the schema provides sufficient clarity to prevent over-sharing while remaining lightweight and usable.

GR illustrated that in non-standardised use cases, EDHs must determine the appropriate data to share based on the consent context, whereas in standardised schemes (e.g. tariff interoperability), this is predefined. This reinforces the need for sufficient granularity in the schema to guide EDHs’ decisions.

A member noted the difficulty of following the design without concrete use cases and recommended linking high-level consumer journeys to the technical schema. GR agreed this would be necessary and acknowledged that the schema may evolve further once tested against end-to-end user journeys and practical scenarios.

GR introduced a potential future capability for CCS to return verified consumer details (e.g. identifiers or MPxNs) back to ATPs following authentication. They asked Members whether such sharing would be beneficial for aligning records across the industry, while acknowledging potential GDPR implications.

Members raised immediate concerns. One member noted that explicit consumer consent would be required for any onward sharing of personal data and questioned who would determine a "single version of the truth," particularly given known challenges with data such as addresses. GR agreed, suggesting that any such functionality would require clear consumer permission and may only provide longer-term benefits in improving data consistency and simplifying processes.

The discussion then returned to schema design, particularly around data type granularity. One member highlighted that, in practice, data types are typically broken down into distinct elements such as import/export, tariff, frequency, and granularity, rather than treated as a single category. They emphasised that these dimensions are already handled as separate fields in existing implementations and should be considered explicitly in the schema.

Another member reinforced that the definition of "type" in the consent schema must align with the underlying data catalogues and services provided by EDHs, including variations in granularity, frequency, and data content. They stressed that schema definitions cannot exist independently of these data structures.

Multiple members cautioned against using consumer-provided data such as email addresses as unique identifiers. They highlighted reliability and consistency issues (e.g. multiple emails per consumer, shared addresses, formatting variations), recommending instead that CCS use system-controlled identifiers (e.g. UUIDs) for data linking and record management.

3. SECURITY PROFILE

GR provided a brief overview of the CCS security profile, explaining that it defines the cryptographic and protocol rules governing the entire ecosystem. This includes how participants authenticate (e.g. mutual TLS), how interactions are secured, and how activity may be monitored and traced across the system.

They noted that the proposed approach is based on FAPI 2.0, which received strong support during the design consultation. However, FAPI 2.0 still requires a number of configuration decisions (around eight key areas), some of which are highly technical and were not covered in detail during the session. GR emphasised that these configuration choices are important but are better reviewed offline by individuals with cybersecurity expertise. Members were invited to review the materials and provide feedback on whether they agree with, or have concerns about, the proposed security configuration for CCS.

4. FORWARD PLAN

The Chair concluded the meeting by outlining upcoming sessions and encouraged ongoing feedback and engagement from all participants.

There is a provisional date for the next Engagement Webinar on 16 July, which will cover a review of RECCo's quarterly plan. The next Topic Solution Group meeting later this month will focus on policy position updates, mandatory introspection, CCS User platform R&R, RBAC and Auth Roles- Data Share Access. Other planned topics for future meetings include SEC & REC alignment, Legal Access Agreement, CCS User requirements, CCS Definition, and ID&V. The Chair advised that Topic Solution Groups will continue on a fortnightly basis, with future topics to be confirmed as soon as possible.

MPW advised that any members with queries are welcome to email consumerconsent@retailenergycode.co.uk, and the RECCo project team will be happy to respond. MPW added that RECCO are looking for early adopters of the CCS solution and asked any interested parties to scan the QR code at the end of the meeting slides.

RETAIL ENERGY CODE

ATTENDEE ORGANISATIONS

Attendee	Representing	Initials	Role
Teresa Thompson	Retail Energy Code Company	TT	Chair
George Roberts	Retail Energy Code Company	GR	Presenter
Matt White	Retail Energy Code Company	MPW	Presenter
Sarah Jones	Retail Energy Code Company	SJ	Presenter
Edma Hyskaj	REC Code Manager	EH	Secretariat

Ofgem	Publicis Sapient
Department for Energy Security and Net Zero	Centrica
Citizens Advice	CGI
Which?	SECCo
Consumer Scotland	SMS Energy
ScottishPower	PayPoint
Elexon	GreenSync
Perse Energy	National Energy System Operator
SSE	N3rgy
E.ON	Flock Associates
East Lothian Housing Association	One Utility Bill
Smart DCC	ESG Global
Raidiam	PRIME
Electralink	ScottishPower
Cadent	Ohme
Stark	Smart Energy GB