Annex D: Technical Design

Centralised versus Decentralised Components

The following paragraphs outline some key boundaries where we see the hybrid design as fitting best for industry and Ofgem's aims with the CCS.

Token Provision and Centralised Ledger

We are proposing to employ a model that includes centralised token provision, alongside a centralised ledger. We recognise that this is a departure from other ecosystems, such as Open Banking, which utilise a decentralised model:

- Open Banking globally tends to operate a decentralised token model, where authorisation for data sharing is handled by banks and financial services organisations themselves. This approach works well, as it avoids the additional cost and duplication of a centralised solution, under Open Banking, banks are required to manage and maintain consent records themselves.
- Additionally, a decentralised ledger/token model avoids the risk of operating a central platform which becomes a single point of failure and therefore necessitates high availability, responsivity, and BCDR plans. We recognise organising consent around a single token mechanism, rather than through distributed implementations, carries risk.
- A decentralised approach may also appear to be a suitable fit for the sharing of personal energy data in the UK, given the existing SEC Other User market, where organisations are accessing personal data, handling consent records independently, and interoperating.

Whilst we did investigate the viability of a decentralised approach, we have determined the following key factors that suggest a centralised model will be best for the UK energy market:

- Ofgem have stated an intention to raise the standard for how consents are managed for energy data sharing. Due to the existence of a range of solutions in the market today for consent management and energy data sharing, we believe there is a significant opportunity to develop a more consistent approach that meets the agreed minimum standard, whilst offering ancillary benefits to those organisations sharing data. An efficient way to do this is to introduce a central solution that can manage this aspect for industry by removing the need for consent management, storage, or tokenisation for existing players.
- Ofgem's requirements clearly state the need for a single source of truth for consent records, necessitating centralised record storage. This approach lends itself to a central token model which can record and store consent records continuously.
- The UK energy market has different characteristics to UK banking. Banks and financial institutions are very advanced in their KYC processes, lending itself to a distributed token model where high baselines of trust are driven by strong regulation and universal implementation. At present, UK energy suppliers are not mandated to have the same KYC processes in place, which would introduce risk if a decentralised token mechanism was introduced, as there would be inherently more risk of consent misappropriation. In addition, there are a range of EDPs which are not energy suppliers and therefore fall outside energy regulations, limiting the regulatory enforcement powers.
- A central model will have notable benefits for auditing and issue arbitration by having an independent register, which should benefit industry as a whole.
- We acknowledge the potential risks of operating a central model (as outlined in bullet three under 6.47), however we believe there are appropriate mitigations. We intend to procure a solution that can provide high availability out of the box and comprehensive BCDR plans. However, while we acknowledge that not all outages can be prevented, we propose contingency plans in Section 5, outlining how the CCS can be operated in the event of an outage.

- A central token model will handle authorisation and authentication in the backend without any actions required by the consumer, except for authentication and consenting. We believe this approach best fits Ofgem's aspirations for a positive consumer experience.

Centralised IDV

Our preferred approach for IDV is to provide a central solution for MMP to ensure robust standards are met, without forcing immediate actions on industry.

- As outlined above, energy supplier KYC processes are not equivalent to those mandated for banking and financial institutions. In banking, strong KYC processes have driven a very high bar for IDV solutions, enabling banks to have high confidence in an account holder's identity and, therefore, rights to services.

- Strong IDV in banking means that banks can be very sure that a consent request, and any subsequent records, were generated by the individual to whom the record relates.

- The lack of an equivalent to this IDV approach in the UK energy market has directed us towards a preference for a centralised IDV solution that can raise the standards in line with Ofgem's goals of enhancing data protection and information security standards in the UK.

- Centralised IDV, alongside a centralised ledger, would provide a single consumer record to mitigate the need for consumers to carry out additional IDV checks each time they engage with a separate service provider requiring consumer consent.

- An alternative would be to mandate suppliers to develop their own IDV processes and handle authentication directly with their customers. This may offer some benefits through reducing reliance of a central solution, but it would lead to more complex implementation plans, with suppliers required to validate all data sharing requests. We believe supplier KYC and IDV upgrades could be a viable solution for beyond MMP, however given project timelines and the significant amount of ongoing, parallel initiatives which industry are having to respond to, we are of the opinion that implementation of a central IDV solution in the short-term is the most suitable answer for industry.

- We are also working to identify the best IDV approaches for consumers through market reviews and consumer research to identify any solutions that can provide the requirements for access to data via consent.

Decentralised Data Sharing

Under the CCS model, it is proposed that energy data is exchanged directly between EDPs and ATPs, outside the CCS technical solution. However, to protect the integrity of the solution and promote trust, DSAs are expected to meet minimum CCS-dictated technical requirements to ensure they meet security and interoperability standards:

- For all cases where an individual's personal energy data is being exchanged within the ecosystem, data sharing will need to meet the minimum security profile requirements outlined above.

- Following our recommended approach to use the FAPI 2.0 protocol, this would only necessitate the exchange of tokens provided by the central solution (cryptographically-bound to the ATP) and the use of mTLS for encryption. This position is driven by the high risk of personal energy data and MPxN data exposure and CCS's key design principles of consumer protection.

- Nevertheless, the boundary for CCS determining technical requirements ends when data has been exchanged between an EDP and an ATP. Therefore, when an ATP is sharing data onto the end consumer or another commercial entity, it is up to that ATP to decide how this data is shared and what controls are in place. From the perspective of CCS, this is where the boundary of the REC stops and the jurisdiction of GDPR begins, since the CCS will govern only how EDPs and ATPs exchange data within the trust framework. ATPs will be at liberty to determine what encryption requirements they deem suitable for that use-case.

- Data sharing between ecosystem participants, however, when data is shared by an EDP to an ATP that would fall within the jurisdiction of the CCS and therefore would necessitate the controls established by the solution. The diagram below outlines this thinking which we welcome your feedback on.

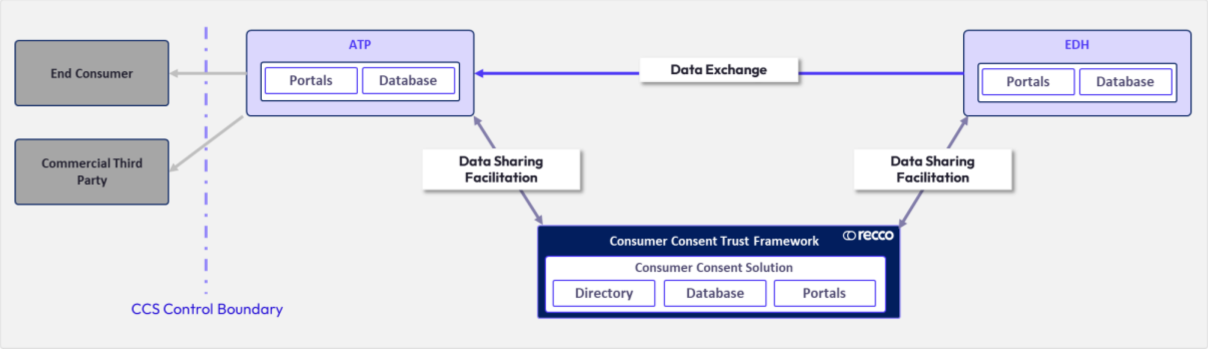

Figure 1: Decentralised data sharing

Standardised Data Formats and Structures

The CCS will define the central Consent Data Schema to be used by each organisation when enabling consent-based data sharing (as outlined in Section 4 of the consultation). This will standardise the data items and formats that are used when consent is requested and stored to ensure interoperability and that all necessary data about a consent is captured.

In essence, this schema is designed to structure the data items to reflect what data is being shared, with whom, and for what duration. Additionally, the central Consent Data Schema must be able to reflect the individual's right to access that data, driven by their occupancy of the address.

In terms of the energy data shared bilaterally between EDPs and ATPs, we do not believe it would be feasible to define a single half-hourly data schema given the range of different use cases that will require the sharing of different data items and data which has different provenance, e.g., half-hourly consumption data shared directly by SEC Other Users accessed via the DCC standard interface may differ from half-hourly consumption data shared by Elexon via the SDR based on standard MHHS message formats. This is something that could be considered further at a later date.

However, to support consistency across the energy data sharing ecosystem, we are proposing that, whilst EDPs will be responsible for defining their own APIs, these should be based on data item definitions reflected within a central data catalogue. In addition, the CCS rules will define the lexicon for how terms of data capture and usage are displayed to the end consumer at the point of consent. This lexicon will be determined according to the CEGs.

We are proposing that the central data catalogue will be delivered through the ongoing use of the Energy Market Data Specification (EMDS) which holds the metadata associated with the majority of existing energy market API messages, including data items used for gas and electricity switching and settlement activities.

Whilst the EMDS is owned and administered by RECCo, the associated governance is open and transparent, with message and data item ownership captured against individual items, enabling the applicable owner to determine the relevant content of the EMDS and direct any required changes. For example, the majority of MHHS messages and data items are owned by the BSC, with change governed through the BSC change process and reflected in the EMDS on direction from Elexon. We anticipate a similar approach for any data items shared via the SDR, noting that these are likely to be based on current MHHS data definitions.

As the CCS scope expands to include wider datasets, we will consider whether a standard Energy Data Schema should be defined on a case-by-case basis. This determination will recognise the difference between many-to-many relationships (where many ATPs are accessing data from many EDPs) and many-to-one relationships (where many ATPs are accessing data from a single EDP). It is proposed that for many-to-one relationships, the

EDP would be responsible for defining the Energy Data Schema, given their overall ownership of the data sharing arrangement (for example, with Elexon's SDR).

For many-to-many scenarios, standardisation is important to foster efficiency in data sharing across industry by minimising effort for data integration through interoperability. It may therefore be helpful to define a single governed Energy Data Schema. This is reflected through the parallel activity being progressed to define a standard API for the sharing of tariff pricing data.

Further work is required as part of the lower-level design activities to determine to what extent standardisation is necessary, as well as to consider the inclusion of non-SDR data items within the EMDS to determine whether this is possible within the MMP timescales.

Centralised Directory and Registry

The CCS will provide a central Directory to enable discovery of all ecosystem participants, as well as necessary machine-readable information to facilitate integration. This central list forms the basis of the Trust Framework and will include a Registry of accessible datasets.

The CCS design aims to foster open data discovery through API-first design and interoperability. However, we recognise the current existence of direct, bilateral agreements in place throughout the market. Consequently, we are working to ensure the CCS ecosystem can enable public and potentially commercially-private data discovery, although the precise technical systems to enable this have yet to be defined.

The CCS Directory and Registry are not being designed to prevent application of commercial arrangements within them. Organisations will be able to establish their own DSAs within the ecosystem should they wish, however, this would require configuration. Any bespoke data sharing arrangement will be reviewed centrally and authorised by RECCo in accordance with the criteria set out within Section 8 of the consultation.

Data Storage

The CCS design does not include a requirement for energy data to be held within the CCS itself e.g., there is no centralised data lake. However, consent records will be stored centrally within the CCS to support the central portal functionality, allowing consumers to view and manage all of their consents in one place. Additionally, a central database will support RECCo assurance and arbitration activities required for the ongoing management of the ecosystem.

CCS Users will still be responsible for the storage and management of their own consent records according to GDPR. The CCS and CCS Users will communicate via webhooks to ensure changes to consent records are updated universally and immediately across the ecosystem.

Central User Interfaces

As outlined above, the CCS will provide a central consumer-facing portal. However, this does not prevent ATPs from maintaining and owning their direct relationships with consumers for the services they provide. Should a consumer wish only to use one ATP for all of their services, they will be able to continue doing so by interacting solely with that provider. The central solution would be updated with information via webhook notifications from the ATP into the central solution.

The CCS will also provide a central administration portal to enable self-serve functionality for CCS Users. This may include services such as certificate management and rotation, API management, and dataset management, with the aim of enabling ATPs and EDPs to manage their involvement within the ecosystem. The aim is to provide a platform that provides the right tools for organisations to self-manage and optimise their involvement in the ecosystem.

Summary

We believe our approach to these nuanced aspects of the CCS provides an effective balance between the centralised elements driving consistency and minimising fragmentation and duplication, and the decentralised elements that recognise the diversity in the data sets, use cases, and EDP governance included within the scope of the CCS.

As part of our ongoing engagement with prospective EDPs, we will further develop lower-level boundaries between CCS responsibilities and elements that EDPs will be responsible for. As an example, we expect EDPs to define the individual data sharing APIs that ATPs will use to access energy data, whilst RECCo will define the consent sharing APIs that ATPs and EDPs will use to exchange consent tokens and validate that active consent is in place. However, further consideration regarding the applicable standards is required, particularly the mechanism for managing security credentials to ensure the overall data sharing ecosystem is designed to minimise friction and duplication for ATPs.

Non-Functional Requirements and Proposed Monitoring

This section of the annex outlines our proposed approach to monitoring and summarises key non-functional requirements.

The CCS will include automated monitoring capabilities covering all components of the service. As a minimum, this will include:

- CCS system availability and resilience;
- API performance, including latency, timeouts, and error rates;
- volumes of consent creation, updates, withdrawals, queries, and revocations;
- IDV success and failure rates;
- address and MPxN matching success rates;
- drop-off rates at key stages of the consumer journey, consent journey durations;
- trend analysis of failed, repeated, or suspicious consent attempts;
- changes in behaviour patterns that may indicate misuse or security concerns;
- tracking of CCS-raised queries against subsequent system actions;
- monitoring of overall demand and capacity (e.g., throughput, peak utilisation) to identify bottlenecks and inform scaling;
- monitoring of data governance, management, and retention; and
- monitoring of disaster-recovery performance.

These monitoring components will enable RECCo to identify issues such as:

- accessibility barriers causing consumers to drop out;
- identity matching failures that may signal incorrect data inputs or system defects;
- ATP behaviour inconsistent with expected consent management practices;
- potential security risks (e.g., repeated unsuccessful consent attempts from the same ATP);
- systemic errors in the identity or address matching process;
- DSAs producing abnormal patterns of access;

- CCS performance degradation due to capacity constraints or high traffic loads;
- CCS failure to meet latency, availability, or accuracy thresholds;
- abnormal error patterns or repeated retries indicating possible system defects;
- failures in CCS backup, data retention, or audit-trail creation processes; and
- failure to recover the CCS within the expected time or without losing more data than planned after an incident.

CCS Users will be monitored to ensure they comply with their obligations under the REC and maintain consumer trust. Monitoring will focus on adherence to required processes, behaviours, and outcomes, and will not constitute validation of the legal sufficiency or appropriateness of individual consents, which remain the responsibility of the organisation relying on that consent. Monitoring will include:

- adherence to CEGs (e.g., clarity, transparency, accessibility);
- timeliness of responding to consumer queries and queried consent investigations;
- non-compliance with the ongoing accreditation requirements described in Section 8;
- alignment of ATP internal consent records with CCS records, including reconciliation outcomes; and
- timely use of the CCS portal to update, correct, or revoke consents where required.

Where patterns indicate potential non-compliance, behavioural risks, or consumer harm, RECCo may apply targeted performance assurance techniques in accordance with the REC PAF, (e.g., audits, information requests, or improvement plans).

Monitoring of CCS Users will focus on compliance, behaviour, and consumer outcomes rather than system performance. Potential reporting areas include:

- periodic reconciliation outcomes (CCS vs ATP records);
- consumer query analytics (volumes, trends, outcomes, root-cause themes);
- material data protection issues identified through CCS-related activities;
- voluntary self-reported improvements or incidents affecting consumer journeys; and
- correct implementation of the retry and error-handling behaviours defined in the API Technical Specification.

The following table summarises the key NFRs that we intend to monitor and review as part of ongoing assurance activities:

Non-Functional Requirements			
NFR Area	Description and Value Definition	Metrics	CCS or CCS User
Responsivity	• Indicates how long requests take to process, a key driver for consumer satisfaction. • Low responsivity could also indicate performance issues helping RECCo understand bottlenecks or particular use-cases that are driving demand for the service, enabling proactive service optimisation.	• Response Latency	• Transparent response times to be defined for the CCS itself. • Response times for the provision of data from EDPs will be driven by the services they offer and agreements with their users.
Availability	• Determines the amount of time the service is running and available, crucial to mitigating the risk of consumer drop off.	• Percentage Uptime	• Important to have clear availability requirements for the CCS itself. • Availability requirements for ATPS and EDPs will be driven by the services they offer and agreements with their users. In further design work, RECCo may determine a minimum uptime for the ecosystem to ensure potential benefits are realised.
Demand and Capacity	• Provides insight into the number of requests for the service and the volume of data being exchanged. Understanding this can illustrate the value of service, its usefulness for industry, and inform future planning for optimisation or scaling.	• Traffic (no. requests per minute)	• Maximum demand volumes to be defined for the CCS itself to reflect the expected usage patterns and ensure sufficient capacity for delivery without impacting message latency. • Maximum demand volumes for EDPs will be driven by the services they offer and agreements with their users.
Data Retention	• Ensures an audit trail is available to support issue and dispute resolution and mitigates the risk of non-compliance with GDPR. Expectations to also be set regarding daily backups and retention of	• Timescale for holding data	• Data retention periods to be defined for the CCS itself and CCS Users. Requirements will include the retention of consent data, audit trails of interactions with the CCS, and the retention of consumer data received by ATPs.

	information to support disaster recovery.		
Accuracy	• Identifying success and error rates indicates the true end-to-end performance in fulfilling the objectives of the CCS. This data can also be used to help with proactive monitoring to minimise downtime and optimise the service.	• Errors • Error codes • Error rates • Success rate	• The CCS API Technical Specification will define the mechanism for granting consent and the associated error codes where requests cannot be fulfilled. • The CCS itself and CCS Users will be required to comply with the CCS API Technical Specification, with mechanisms in place to monitor performance.
Attrition Rates	• Measures the consumer engagement throughout the CCS customer journey to understand any areas of friction leading to consumers abandoning the process.	• Drop off rates	• Requirements will be defined for the CCS itself to minimise consumer drop off, with iterative service improvements focusing on this element.
Retry Strategy	• A clearly defined, automated, retry strategy provides confidence to users where data / consent requests are not fulfilled. This will include a buffer period after which the user suspend retries.	• N/A	• The CCS itself and CCS Users should all be required to implement the defined re-try strategy.
Query Resolution SLAs	• The ability for consumers to query consent data will be built into the CCS. • Defined SLAs for managing queries / issues will mitigate the risk of negative consumer experience.	• Query response • Query resolution	• There will be a variety of issues and potential errors that may arise through interactions with the CCS. These will be mapped into a set of error resolution paths with clearly defined steps. It is anticipated that defined SLAs will be applicable to the CCS and CCS Users.
RTO and RPO	• Clear requirements to reflect the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) to ensure that where outages occur data loss is minimal and recovery activities achievable.	• Downtime • Data Loss	• It is important to have clear RTO and RPO requirements for the CCS itself. • Requirements for ATPs and EDPs will be driven by the services they offer and agreements with their users.