



Data Standards Body

Non-functional Requirements (NFR) Consultative Group

Minutes of the Meeting

Date: Wednesday 14 August 2024

Location: Held remotely, via MS Teams

Time: 14:00 to 16:00

Meeting: Meeting # 7

Attendees

Participant Members

Mark Verstege, Chair
John Adshead, AEMO
Brenda Ashcroft, AGL
Jon Denley, Basiq

Dhananjay Gourshettiwar, Westpac
Harish Krishnamurthy, ANZ
Thomas Lu, Origin Energy
Julian Luton, CBA

Observers

Elizabeth Arnold, DSB
Nils Berge, DSB

Terri McLachlan, DSB
Hemang Rathod, DSB

Apologies

Jim Basey, Basiq
Michael Lin, NAB

Mark Wallis, Skript



Chair Introduction

Mark Verstege, the Chair of the Non-functional Requirement Consultative Group (NFR CG) welcomed everyone to the meeting and acknowledged the traditional owners of the lands.

The Chair noted that Jim Basey (Basiq) was an apology for this meeting. He also welcomed Thomas Lu to the group as Origin Energy's representative and Brendan Ashcroft who has replaced Andrew Ferris as AGL's representative.

Minutes

The Chair thanked the group for their comments on the minutes from 17 July 2024 meeting. The Minutes were formally accepted and will be made publicly available on the Consumer Data Standards website.

Action Items

The Chair provided an update on Action Items as follows:

- Xero has provided some high-level metrics and data collection patterns with further discussions planned. Xero may attend the next meeting
- Invite energy representative to the group - completed
- DSB to complete analysis of solution components back to themes and issues – ongoing
- DSB seeking feedback on the retro – will be addressed at meeting
- DSB to create a survey inviting ADRs to comment on data collection patterns and mapping use cases – ongoing
- Data holders to provide feedback on cost considerations and implementation feasibility of different TPS thresholds – ongoing

The group highlighted the significant effort and cost involved in scaling infrastructure to support higher TPS. The DSB responded saying that it's their assumption to not increase TPS but maximising the use of the TPS currently provisioned.

Terms of Reference

The Chair thanked the group for their comments and the updated Terms of Reference were formally accepted and will be made publicly available on the Consumer Data Standards website.

Review of Retrospective

The DSB conducted a retrospective focused on evaluating the group's progress and identifying areas for improvement. Key points included:

- Simplification based on products and data sharing eligibility rules was suggested to reduce the load on infrastructure and processing. This involved reducing the number of products in scope and simplifying the rules for data sharing eligibility to minimise the number of API calls and checks required for each data sharing request
- We are open ended. What is the product of the consultative group - how do we know when we get there?



- Consider if any new rules changes should be taken into account - or future standards changes (e.g. RAR and implications for data clusters and API calls).
- Acknowledging the need to identify key issues and work on them in a time-boxed approach.
- The importance of continuing discussions on both ADR and data holder sides, focusing on efficient usage and removing duplicate requests.
- To create GitHub issues for broader community feedback on proposals.
- Consider organising in-person meetings for deeper engagement.
- Re-evaluating the CDR's direction, considering the high volume of repeat transactions and the potential for more efficient data sharing patterns.

The DSB suggested that the next steps could be to start drafting a report focussing on facts, findings and recommendations based on the discussions and analysis conducted and start to consult on change.

Update on metrics data analysis

The DSB discussed the metrics data analysis, focusing on capacity planning based on the number of active arrangements and the impact of large accounting platforms like Xero on capacity. It was noted that Xero's data collection occurs within a specific window each day, which does not maximise the 24-hour period, suggesting a need for better guidance to smooth TPS load throughout the day.

The discussion also highlighted the importance of making the system more efficient within current TPS thresholds to accommodate customer growth without additional infrastructure scaling.

The DSB recapped the key problems they are solving:

- Improving the sharing of large quantities of data: This involves finding more efficient ways to share data that doesn't change frequently or is historical, ensuring data recipients can access the data they need without unnecessary API calls.
- Improving the sharing of low velocity data: focusing on how to share data that changes infrequently, like standing data in energy or historical banking transactions, in a way that reduces the need for repeated requests.
- Improving the sharing of data that has changed: Identifying methods to allow data recipients to easily determine what data has changed since their last request, to avoid fetching large volumes of data where only a small portion may have changed.

The DSB noted that the current state of the meeting discussions focused on improving the efficiency and capacity of the CDR ecosystem. Key points included:

- Improving Data Sharing: Exploring ways to share large quantities of data more efficiently, including low velocity data and data that has changed.
- Minimising API Calls: Considering bulk arrangement APIs and last modified queries to reduce the number of API calls needed for data reconciliation and updates.
- Increasing Ecosystem Capacity: Aiming to add more capacity within the current system to accommodate growth in the number of data consumers without increasing infrastructure costs.
- Quality of Service: Discussing potential for better quality of service and fault tolerance in data sharing processes.



- Infrastructure Costs and Complexity: Seeking to reduce these for both data recipients and data holders through more efficient data sharing mechanisms.

These discussions aim to make the CDR system leaner, more efficient, and capable of supporting more data consumers while maintaining or reducing infrastructure costs and complexities.

The DSB noted the target state aims to achieve a more efficient and effective data sharing process within the current TPS thresholds, maximising the use of existing infrastructure, and potentially reducing costs by minimising unnecessary data requests and optimising the sharing of data based on its velocity and change frequency.

The DSB noted the outcomes they are trying to achieve are:

- Reduce volume of API calls
- Limit number of API calls to get essential data
- Allow for TPS smoothing and QoS levels
- Reduce infrastructure costs for DHs
- Reduce redundant API request workloads
- Reduction in net TPS
- Reduce processing complexity for ADRs
- Improved data consistency: ensure data is not missed or changes go unreconciled
- Service More customers per software product (increased capacity)

The DSB sought feedback from the group on the following topics:

Option 1: Bulk arrangement approach

The bulk arrangement approach aims to minimise the number of API calls and arrangements a data recipient needs to process. It's designed to solve for intraday and interday data patterns but not for historical data collection. This approach could lead to a reduction in TPS, infrastructure costs, and processing complexity for data recipients. However, it would require the introduction of a new API and maintaining state against arrangements, which could add complexity. The approach seeks to inform data recipients about which arrangements have changed, potentially reducing the volume of data they need to process

Feedback included the potential of a bulk arrangement API to minimise API calls by allowing data recipients to identify which arrangements have changed. Concerns were raised about the feasibility and implementation complexity, especially regarding tracking changes across multiple scopes and accounts.

Option 2: Last Modified query/responses fields

The "last modified" query response field was discussed as a potential solution to improve the efficiency of data sharing, particularly for intraday and interday data patterns. This approach would limit the result set to only include transactions that have changed since the last modified date specified by the data recipient. It would not reduce the number of customers or arrangements iterated but could provide efficiency gains by reducing the response rate and potentially decreasing the number of API calls needed for reconciliation. This solution is seen as compelling, especially for banking transactions and possibly for energy usage data, considering rate quality changes.



Feedback included practicality of this approach for energy data and concerns about the implementation complexities and the dependency on vendor systems. The discussion underscored the need for solutions that can efficiently indicate data changes without extensive system overhauls.

It was noted that transaction detail retrieval could impact the load and suggested that if data recipients could avoid retrieving transaction details for already reconciled transactions, it could service more calls with the same infrastructure.

Data Collection Patterns for Minimising API calls

The DSB introduced a trial for asynchronous data collection as a potential solution to improve the efficiency of data sharing, particularly for banking transactions and energy usage data. This approach could help reduce the volume of API calls and allow for smoother TPS across the day. However, the feasibility and implementation details remain to be explored further.

The objective is to determine the feasibility of the "429 Retry-After" pattern in handling large data requests, potentially reducing the processing burden on data holders and improving response times for data recipients. They proposed to timebox the trial for a period of one month.

The DSB outlined the measurement criteria which will help assess the effectiveness of the asynchronous pattern and guide decisions on its broader application.

Meeting Schedule

The Chair noted that the next meeting is scheduled for Wednesday 11 September 2024.

Any Other Business

No further business was raised.

Closing and Next Steps

The Chair thanked the group for participating for attending the meeting.

The Chair noted that the Action Items/Next steps include:

- Further exploration of asynchronous data collection patterns
- Potential drafting of a report summarising the findings, facts and recommendations which could provide a comprehensive overview of the discussions and outcomes
- DSB to send the design of the asynchronous patterns to the group for review

The meeting closed at 12:00