



Data Standards Body

Non-functional Requirements (NFR) Consultative Group

Minutes of the Meeting

Date: Thursday 17 July 2024

Location: Held remotely, via MS Teams

Time: 14:00 to 16:00

Meeting: Meeting # 6

Attendees

Participant Members

Mark Verstege, Chair
John Adshead, AEMO
Jon Denley, Basiq
Dhananjay Gourshettiwar, Westpac

Harish Krishnamurthy, ANZ
Julian Luton, CBA
Mark Wallis, Skript

Observers

Elizabeth Arnold, DSB
Nils Berge, DSB
Ruth Boughen, DSB

Terri McLachlan, DSB
Hemang Rathod, DSB

Apologies

Michael Lin, NAB



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 the Data Standards Chair approved the extension of the Consultative Group trial for a further six months with the existing membership. They will also seek to recruit some additional people from the energy sector.

The Chair noted that Michael Lin (NAB) was an apology for this meeting.

Minutes

The Chair thanked the group for their comments on the minutes from 19 June 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 the outstanding Action Items as follows:

- Continued liaison with Xero about attending a future meeting
- Liaising with multiple energy retailers about involvement in group
- Completed further analysis and discussed the primary drivers, issues, and opportunities for the NFRs
- Continuation of data analysis to overlay current TPS for organic traffic with new figures
- Continued discussion on load management and exploring solutions to manage peak and average TPS effectively

Update of traffic forecast

Primary drivers, issues and opportunities

Nils Berge from the DSB gave a summary of the main issues and themes that emerged from the previous meetings and the data analysis on the NFRs. He grouped the issues into five themes:

1. Large number of accounts: how to handle consumers with many accounts, especially in the energy sector, and reduce the number of calls and response times.
2. Efficient requests/bulk data/low velocity data: how to reduce the number of calls for low velocity data, notify changes, filter transactions, and make bulk data requests.
3. Customer present awareness: how to support the secondary data holder scenario for energy and follow the NFRs based on the customer present flag.
4. NFR settings: how to review or refine the NFRs such as TPS, response time, page size, and high traffic period, and whether they are appropriate for different types of data and use cases.
5. Develop/enable efficient request patterns: how to explore and agree on different request patterns, such as intraday, interday, and historical, and how to support them with the standards.



The DSB presented some data analysis on the active authorisations, peak and average TPS, and recipient count for eight banks and four energy data holders, and showed some trends and variations across them. They linked the data analysis to the three patterns and suggested some potential solutions for each of them, such as bulk arrangement API, last modified query, event-based notification, and statement API. They asked for feedback and prioritisation from the participants on the themes and the solutions.

One member asked about the difference between the average and the max TPS, and whether we capture the average response time because if the overall response is slower, the TPS will be higher?

The DSB noted this is captured in the metrics but not presented in these charts. They can however look at potentially including for example error accounts reported or rejections overlaid against the chart.

The DSB noted that the upwards trend in the growth of numbers of software product connections and the growth in the number of active authorisations i.e. in the range of 30% to 40% increase over five months in the number of authorisations which indicates the increase usage in the CDR. For TPS tiers, none of the data holders (DHs) are hitting their ceilings within the TPS tiers. When considering peak TPS versus average, there's a huge amount of underutilised capacity and inefficient request patterns.

Data Analysis Update

The DSB provided an update on the use cases and scenarios that the group have been exploring based on the data analysis and the issues and opportunities identified in previous meetings.

They presented three main patterns of data collection: Intraday, Interday and Historical and explained how they relate to different consumer needs and use cases.

Some design limitations include max cap on API invocations per day per software product which creates a cap on number of total daily active customers per software product. OAuth based resource patterns are bound to the authorisation; and active authorisation tiers set an upper limit on the max number of active software products.

They want to solve for ecosystem efficiency and scalability; invocation spreading / TPS smoothing; minimising unnecessary calls; allowing ADRs to discover what data has actually changed; reducing call volumes; and shift problem away from endless infrastructure scaling.

They suggested some potential solutions:

- Discover which consumers have data to collect: options of bulk arrangement API; last modified query; and batch data feed.
- Allow DHs to manage TSP smoothing: options of retry-after pattern; and event-based notification
- Bulk detail data collection: options of bulk transaction detail; statement APIs; and no change.

They asked for feedback and prioritisation from the participants on the patterns and the solutions.

One member emphasised the importance of analysing data collection patterns to determine the distribution of traffic across different use cases. This analysis is essential to prioritise solutions that will effectively address the ecosystem's spiky traffic and ensure that the right issues are being



targeted. It would be good to understand how ADRs are utilising data, which will shed light on the most common patterns and inform the prioritisation of solutions to capture the variety of use cases and the frequency of data requests to better manage traffic within the ecosystem.

One member understands these solutions will help in reducing the number of APIs but they feel the amount of work that's needed from the DH side would be significant.

The DSB noted that the discussion has been based around the idea of the ADR sending through information and not the DH keeping track of it. They also noted that if they moved to an event-based pattern that would put the onus back onto a DH.

One member noted they provide a subscriber service to their retailers where if there is a change in one of their data sets (ones that change every 6-12 months) they send a change notification. They suggest it's easier to do a push from them if there's a subscriber.

Further conversation included considerations on how understanding data collection patterns could influence the Transaction Per Second (TPS) thresholds. By identifying the predominant patterns, there may be opportunities to optimise these thresholds, potentially reducing infrastructure costs for DHs.

The group discussed strategies for maximising efficiency and scalability within the ecosystem by spreading out invocations across the day. This approach aimed to reduce peak TPS spikes and utilise available capacity more effectively.

A number of members expressed a preference for prioritising solutions that enhance request efficiency, such as implementing a mechanism to retrieve a list of accounts with changes since a specific date. This solution is favoured as it promises to improve efficiency without incurring significant costs or requiring extensive architectural changes.

The DSB noted that they will come back with the next level of thinking around the options that were identified and map it back to the themes they are seeking to solve.

The group voted on their preferred options for "Prioritisation" on the Miro board with Theme 2: Efficient requests resulting in 3 votes, Theme 4: Review NFR settings resulting in 2 votes and Theme 1: Large number of accounts resulting in 1 vote.

The group also voted on their preferred options around "Solution components" on the Miro board with Bulk Arrangements API with 2 votes and Last Modified query with 1 vote. Some of the group required more time to consider the options with their teams.

Retrospective

The Chair noted that due to meeting time constraints he would provide members with the retrospective questions for feedback.

Meeting Schedule

The Chair noted that the next meeting is scheduled for Wednesday 14 August 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:

- DSB to complete further analysis of the solution components and mapping them back to the themes and issues
- DSB to send out an email seeking feedback on the retrospective
- DSB to create a survey and invite ADRs to comment on data collection patterns and mapping use cases
- DHs to provide feedback on the cost considerations and implementation feasibility of different TPS thresholds

The meeting closed at 15:59