ESO

Quick Reserve:
Maximum Recovery Period

Justification and analysis

1. Aim

This document follows feedback from market participants and aims to provide further justification behind the 3-minute Maximum Recovery Period requirement for Quick Reserve (QR) and why it should not be increased any further.

2. Context

Recovery Period is the interval between the end of delivery of one instruction to the start of the next. Along with the Time to Full Delivery and Ramping Envelope, these QR parameters are inextricably linked to a system need. Reserve is sterilised during recovery periods, and therefore reducing the recovery period duration means that the reserve can reset and be available for another utilisation sooner. This is linked to the system need to be able to control frequency deviations which occur in short succession. These events happen now, and the trend of future system changes (such as increased volatility and decreasing inertia because of more asynchronous and renewable generation) mean that the ESO view is that pre-fault frequency deviations will continue to occur in quick succession, and it is important to develop a service capable of managing them.

The ESO initially undertook analysis to determine the value of the required Maximum Recovery Period for the Quick Reserve service. The conclusion of this analysis showed that an appropriate Recovery Period would be 1 minute. However, following feedback from early industry engagement events, this requirement was later relaxed to 3 minutes to find the best balance between system security and maximising market participation. Data was analysed between 2014 – 2021 at 1-second resolution. Reference to the previously published analysis via our webinars can be found on the ESO website here.

3. Industry Feedback

Feedback from initial engagement with industry has been mostly positive and in support of the QR Technical and Procurement designs. However, a major topic of discussion has been the Maximum Recovery Period, currently set at 3 minutes, which has received mixed feedback. Most providers agreed with the current requirement, whilst others have tried to argue a case for both shorter and longer Recovery Periods.

In particular, the most flexible providers have argued that the recovery period should be reduced in order that the value of their flexibility to the ESO can be appropriately rewarded, and that less flexible assets should participate in less flexible, lower value services. Conversely, providers with less flexible technology have argued that increasing the period up to around 10 minutes would allow increased industry participation in the service, which would lower costs through increased competition.

The QR service is fuel agnostic which offers many positives – most notably an increased total market capacity, leading to competitive pricing behaviour. However, it does also provide challenges, with each technology type offering different capabilities and limitations. Satisfying all their concerns, which are sometimes opposite, to the needs of the system is therefore a challenging task.

4. Further Analysis

With this background, the ESO has performed further analysis to test the recommendation of a 3-minute Maximum Recovery Period in four key areas:

Available Capacity

Figure 1 below shows available Capacity (MW) plotted against Recovery Period (mins) for assets including BM Battery, Pump Storage and Gas providers.

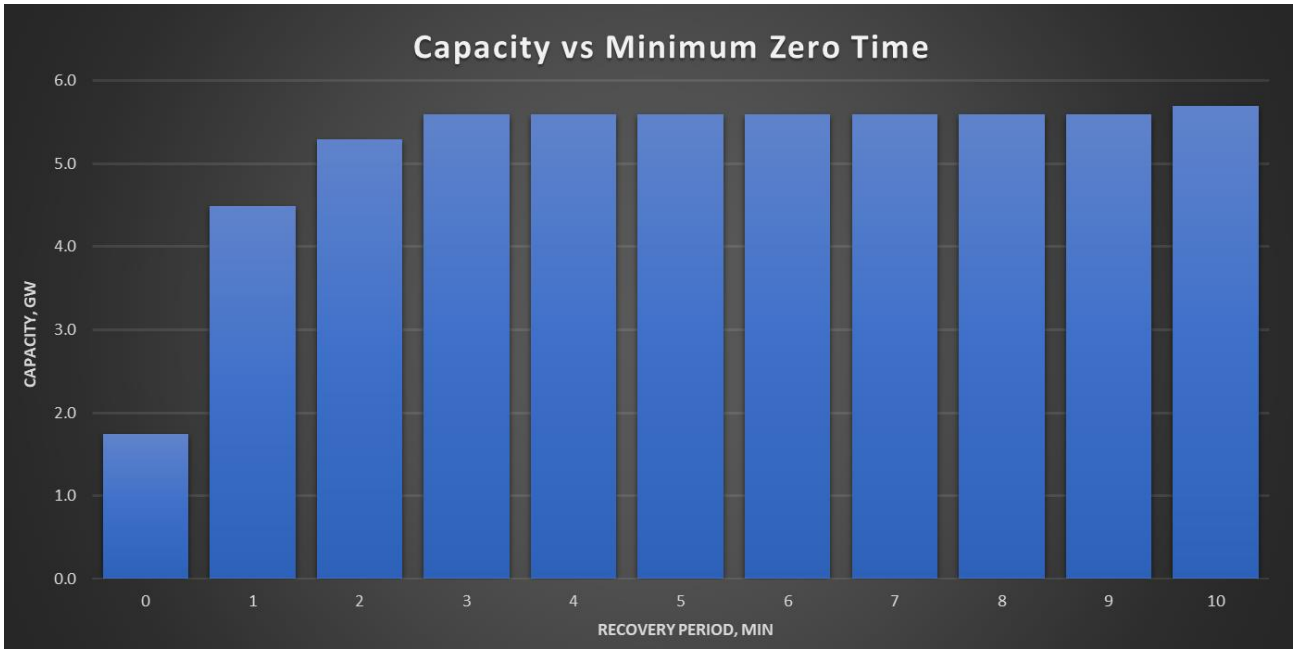

Figure 1 - Available Capacity by Recovery Period: Batteries, Pump Storage & Gas.

Battery and Pump Storage providers are considered key participants for the future Quick Reserve service due to their well-known fast delivery and Recovery Periods, with almost 5GW of potential capacity against a typical QR requirement of ~300MW.

Also included in the above analysis are Combined Cycle Gas Turbines, with around 650MW of capacity. These values are determined through the assumption that these units are already generating (non-zero baseline), have a minimum 1-minute ramping capability of 10MW/min, and have a 1-minute Recovery Period.

The graph shows that there are significant increases in capacity up to 3 minutes, where it then plateaus up to 10 minutes, indicating that there are no benefits to increasing the Recovery Period any further.

Analysis on the wider market is somewhat more difficult, as other types of generation/demand often have variable dynamic parameters and may rely on participation when only at non-zero baselines, along with many other factors.

Frequency Event Coverage

Figure 2 shows Coverage plotted against Recovery Period. Coverage is defined here as the percentage of 0.1Hz deviations a theoretical unit could respond to given a certain Recovery Period. The data only considers time between 0.1Hz deviations ≥1 minute. Anything below 1 minute is assumed to be the target for Response as an automated activation.

It can be seen that units with a Recovery Period of up to 1 minute can theoretically respond to 100% of deviations that fall within the remit of Reserve. If a unit within a contracted period had a 3-minute Recovery Period then it could theoretically respond to ~50% of deviations. As would be expected, this figure drops significantly to only 15% at a 10-minute Recovery period. This graph therefore highlights the importance of a short Recovery period.

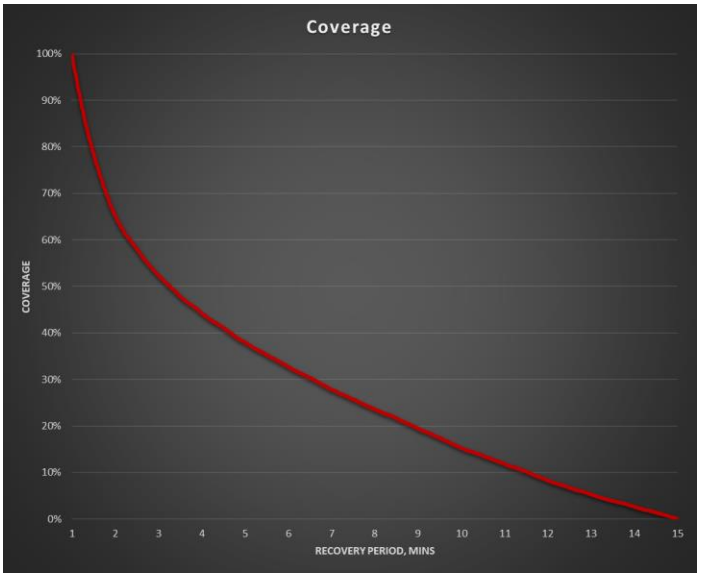

Figure 2 - Coverage Plot, ≥1min deviations

Flexibility

A longer recovery time increases the risk that a unit will be unavailable when it is needed, and thus increases the volume of the service which must be procured. Increasing the Recovery Period is therefore only acceptable if we anticipate a commensurate drop in QR prices.

The ESO has undertaken analysis to assess whether assets would be utilised at the extremes of the flexibility available. Whilst there is limited data to draw on before Quick Reserve as a service is implemented, we looked at instructions to an asset base that has the required level of flexibility. This gave around 4% of instructions that were sent to the same unit with a recovery period of 3 minutes or less, whereas over 20% of instructions were sent with a recovery period of 10 minutes or less. It should be noted that this likely significantly underestimates the instructions sent at shorter recovery periods as these are more likely to be sent to different units due to a lack of flexibility. However this still shows the significant difference that increasing the recovery period to ten minutes would have, based on historical dispatch data.

To try and find an optimum balance between the competing aims of minimising the recovery period for ESO needs, and maximising the recovery period for some industry participants, we have taken the percentage of instructions excluded by each recovery period, and then divided that by the length of the recovery period to get an average impact per minute of recovery period. This therefore will be lower if beneficial to the ESO (by excluding fewer instructions) and lower if beneficial to the industry by having a longer recovery period. In the graph it can be seen that 3 minutes is a minimum, with the next best option being a reduction back to 2 minutes.

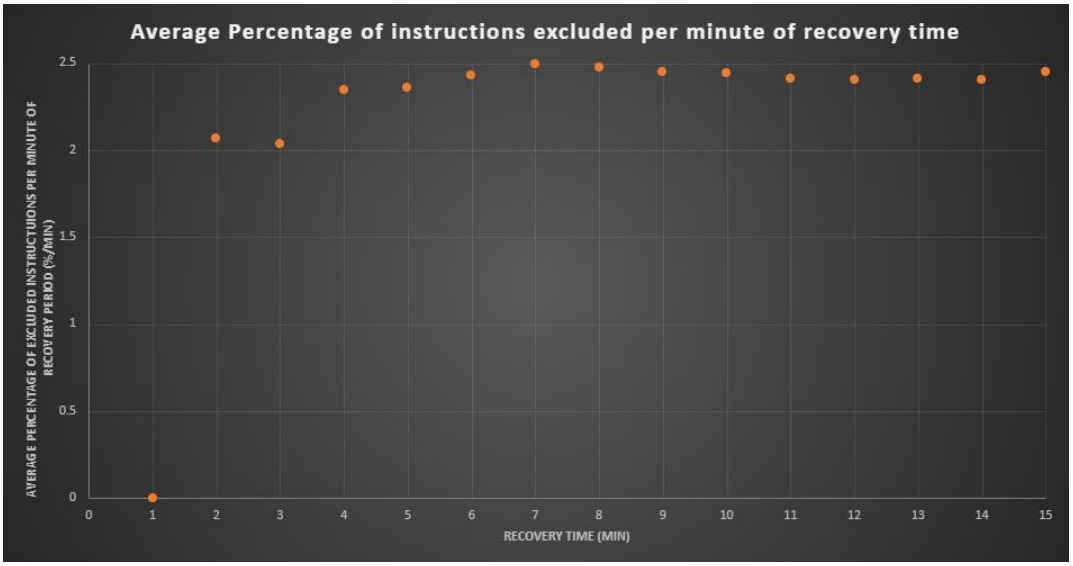

Figure 3 - Combining objectives to find the optimum recovery period

It should be noted that allowing a wide range of acceptable Recovery Periods makes it much more difficult for the ESO to value more flexible units. To accommodate different levels of flexibility - rather than have a single service to accommodate all Reserve – Reserve is split into services that can standardise the terms for similar assets qualities. Assets with ~10-minute Recovery Periods are more aligned with the Slow Reserve service than Quick Reserve, even if they can deliver some aspects of the faster service.

Availability

The risk profiles below are a proxy for availability and, in the case of a single participating unit, is inversely proportional to the number of activations. They show the fraction of time participating units could spend in their Recovery Period relative to the contracted Settlement Period for typical 3-, 4- and 5-minute activation periods. Increasing the number of units within the service does logically reduce this risk profile, as can be seen between *[Left]* and *[Right]* below (note change in y-axis).

For the ESO to see the QR service as reliable, it would need to have an acceptably low risk factor, such that we could guarantee a QR unit will be available when it is needed. The risk of unavailability, when considering a typical participation of 3 units, associated with a 10-minute Recovery Period is significantly higher than for ≤3, prohibitively so at between 4-6x higher. Without this, there is a risk that continued use of other optional services is used to meet this requirement, if required to meet the system need.

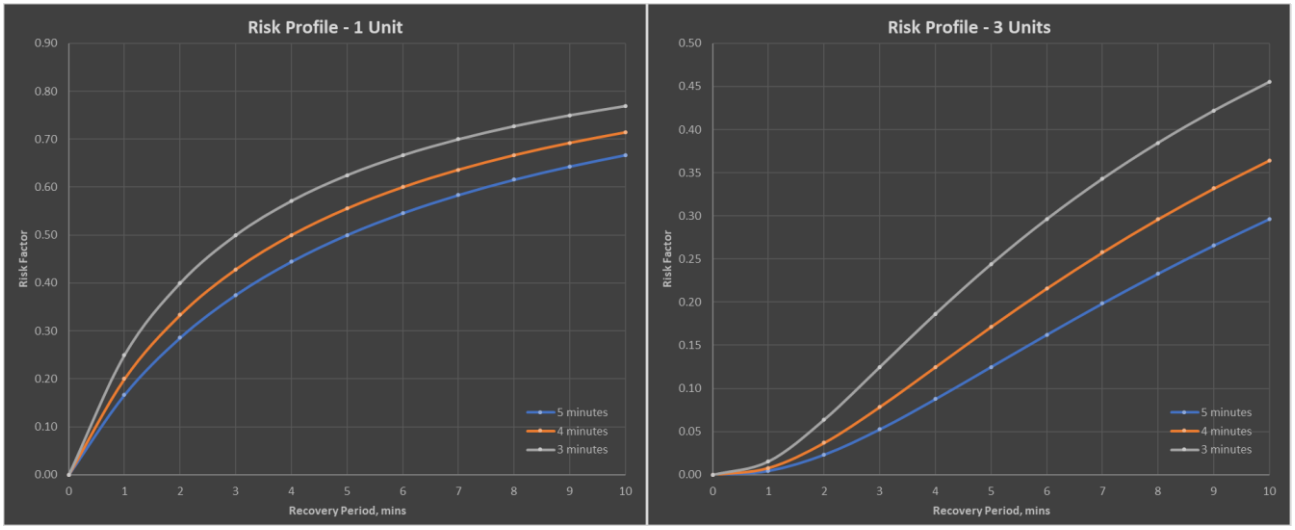

Figure 4 - Risk Profiles [1 and 3 participating units, respectively]

5. Mitigations for increased recovery period

It is possible to mitigate the risks mentioned in section 4 through increasing the procurement volume as, if an asset is sterilised during a recovery period, then having another separate asset ready to dispatch would alleviate the issue. However, the impact on volumes required is significant, as

Figure 5 below shows the required additional volume needed to meet the QR requirement for a given settlement period. The analysis is based off our original 1-minute Recovery Period and considers a typical 5-minute activation period.

The maximum Recovery Period has already been increased from 1 to 3 minutes, which translates to an increased procurement volume of +33%. In an extreme case, raising this to 10 minutes would add 150%, resulting in ~750MW as a typical requirement when compared with an expected 300MW typical requirement. This not only significantly increases procurement costs, but also increases the risk that the necessary required volumes cannot be procured on tight days. Although there is an increased procurement volume for the move to 3 minutes, we believe that diversifying the market, especially away from a single technology type, will bring benefits through competition and co-optimisation by allowing a greater transferability of market volume between different markets. (For example, if only batteries could provide Response and Reserve, the procurement cost would have significant exposure to the marginal cost of battery utilisation, which would be a risk if that spiked above general market costs). However, we do not believe that the additional procurement cost above 3 minutes can be realistically offset by an increase in market capacity leading to greater competition in what we expect to be a liquid and competitive market.

Similar to Balancing Reserve, QR is another new service aimed at reducing total balancing costs. Current fast acting ancillary services like Optional Fast Reserve and pump storage are expensive and uncompetitive, and so QR should, by way of securing competitive volume at day-ahead, reduce these costs significantly. It is therefore pertinent to QRs use-case that we can attain as much value from the service as possible.

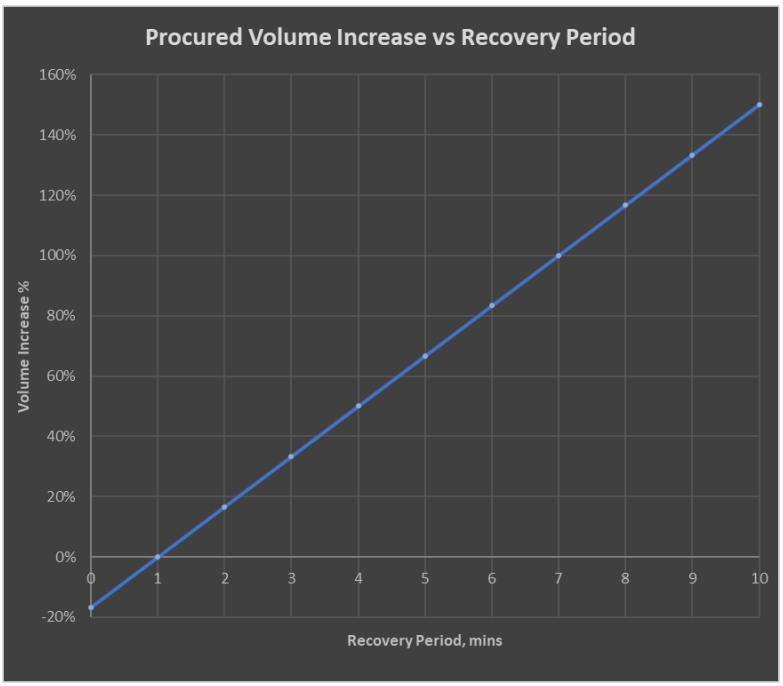

Figure 5 - Procured volume increase vs Recovery Period

6. Conclusions

The ESO previously conducted extensive internal analysis on historic frequency data and concluded that 1 minute was an appropriate requirement. Engagement with industry highlighted a need to extend this value to 3-minutes, which the ESO felt was appropriate to best balance system security against market participation.

In summary, the general effects of increasing the Recovery Period are:

- No additional benefit to the total MW capacity, even when significantly relaxed to 10 minutes or greater.
- Reduces effectiveness of the service by reducing flexibility.
- On tighter days, reduces likelihood that requirement will be filled.
- Increases risk and lowers availability, therefore lowering confidence in its use.
- Procurement volumes would need to increase, which would increase costs for a service which would have a lower technical capability and therefore lower value to the ESO.

Analysis in this paper considers under-frequency events, i.e. positive Reserve actions, and does not consider capacity from Wind. We do not expect wind to participate in positive Reserve with current subsidy arrangements, however they can participate in the negative service, which would lead to significant volumes (GWs) available with 1 minute recovery, therefore making the case even harder to make for increasing the Recovery Period.

There is also the option of reducing the Recovery Period back down to 1 minute. However, this risks excluding market participants and therefore increasing procurement costs through relative scarcity. It is recommended, however, that this is kept under review once the market is established to determine whether there is sufficient liquidity to reduce this parameter in order to better align with system requirements.

It is therefore our recommendation, given all the evidence shown in this and previous justification material (as linked in Section 2), that the Recovery Period remains at a maximum of 3 minutes.

7. Appendix

Consideration for a 10-minute Recovery Period – Best case

Figure 6 below illustrates the "relative volume" available against Recovery period. This relative (or weighted) volume is the same volume data shown in Figure 1 weighted by the percentage values shown in Figure 5. This is to illustrate the potential value of offered capacity relative to its Recovery Period. The data also includes the maximum theoretical volume from additional provider assets - at 3GW as a best case - capable of recovery in 10 minutes (otherwise capable of meeting the technical requirements for Quick Reserve).

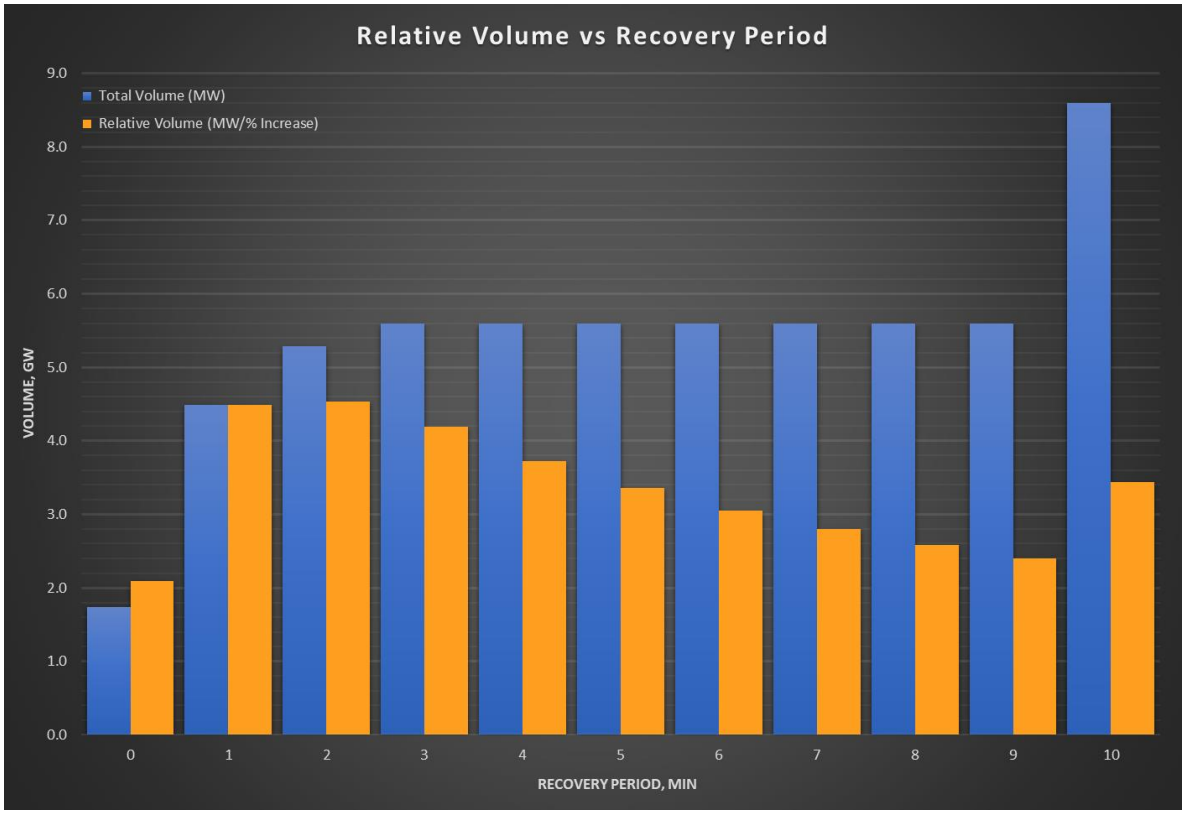

Figure 6 - Weighted Volume vs Recovery Period

For there to be any reasonable benefit to extending the Recovery Period out to 10 minutes, the relative volume available should at least match the highest point on the graph- which can be observed as 4.5GW at 2 minutes. However, at 10 minutes the relative volume is only 3.4 GW, 1.1GW less than the peak. To justify increasing the Recovery Period to such extreme lengths, there would need to be a further 3GW+ of available volume at 10 minutes for the relative volume to equal the same 4.5GWs seen at 2 minutes. It should also be noted that future growth in assets meeting the 3-minute period (such as more installed battery capacity) has not been shown here and would increase the case further for the shorter recovery period.

Figure 6 shows that increasing the Recovery Period, and consequently procurement costs, can therefore not be justified with the above "best case" and, in parallel with Figure 3, makes the case for a reduced Recovery Period of 2 minutes.