

May 2026

Boundary Flow Smoothing Feasibility Report

See separate document for appendices

Contents

1. Executive Summary	4
Context and Objective	5
Approach	5
Key findings	6
Assumptions and Limitations	7
Technology Suitability and Opportunity Cost	7
Conclusions	8
2. Introduction	9
Problem Context	10
Overview of Proposed Solution	11
Boundaries of Interest	12
Report Structure	14
3. Research and Theoretical Feasibility	15
Section Overview	16
International Comparison Projects	16
Historical Boundary Flow Behaviour	17
10-Minute Averages – $\Delta(P90)$	20
10-Minute Averages – Evolution of the Mean	21
4. Cost Benefit Analysis Approach	25
Introduction	26
Routes to Smoothing Benefits	27
Improved Asset Reliability	27
Reduced Requirement for Constraint Actions	28
The Impact of Nested Boundaries	30
Method for Quantifying Smoothing Costs and Benefits	31
Identification of Trial Smoothing Algorithms (1)	32
First Difference Analysis & Reduction in Implied Safety Margin (2)	32
Nested Boundaries Analysis (3)	34
BOA Cost-Saving Analysis (4)	37

Cost Benefit Analysis..... 38

5. Cost Benefit Analysis Results..... 41

6. Discussion and Conclusions 53

Overview of Findings and Sensitivity Assessment..... 54

Potential for Constraint Cost Reduction..... 56

 Interpretation of Estimated BOA Reductions 56

 Energy and Power Requirements for Smoothing Operation 57

Comparison of Costs and Potential Benefits 60

Technology Suitability and Opportunity Cost..... 60

 Technology Suitability 60

 Opportunity Cost - BESS..... 63

Summary Conclusions 65

Suggested Further Work 66

 Operational and Decision-Making Refinements..... 66

 Targeted Network and Boundary Modelling 66


 Asset Behaviour and Real-World Deliverability..... 67

 Market, Commercial and Service Design Development..... 67

 Data Quality and Scenario Expansion 67


7. Glossary..... 69

8. References71



1. Executive Summary





Context and Objective

Power transfers across transmission boundaries on the Great Britain (GB) electricity network are increasingly constrained by **thermal, stability and voltage limits**. Volatility in the flow of energy across transmission boundaries adds to the complexity of managing these constraints and can increase the number and magnitude of interventions (instructing assets to turn down/up) required by NESO's control room to ensure a secure system.

In 2024 alone, thermal-related constraint actions cost approximately £1.4 billion. The majority of these costs resulted from turning down low-carbon generation in areas where it was not possible to export all energy generated to areas of demand, whilst compensating for any resulting shortfall in energy, by generating replacement electricity in unconstrained areas closer to demand (often using high-carbon fuels).

While transmission reinforcement remains the primary long-term solution to resolving many thermal constraints, it will take time to deliver. In the interim, smarter operation of the existing network could help reduce these challenges and the associated costs.


This report assesses a potential measure that could support more efficient use of the existing transmission infrastructure: **a smoothing service** that dispatches flexible assets to reduce short-timescale volatility in power flows across transmission boundaries. In principle, this is achieved by reducing short-lived peaks and troughs around average flow levels, allowing the system to operate closer to its limits without increasing operational risk. This approach will be tested to determine whether it enables increased power flows across transmission boundaries while reducing the need for Balancing Mechanism (BM) adjustments (via bids and offers). Results will be assessed against historical boundary flow behaviour to ensure no increase in the risk of exceeding capacity limits.

Frazer-Nash Consultancy (FNC) were appointed to conduct an independent assessment of this concept, funded by NESO using the **Network Innovation Allowance** (NIA is a set allowance that each network receives to administer as part of their network price control) - **resulting in FNC producing this report** (and separately published appendices).

Periodic updates have been presented during the project at the Constraints Collaboration Project (CCP) industry webinars.

Approach

The analysis uses high resolution historical data on boundary flows and constraints from 1 January 2024 to 31 March 2025, covering a **set of nested boundaries** (see [glossary](#)) **in Scotland and northern England**. Flow data is analysed in 10-minute windows, with the change in mean flow between consecutive intervals used to characterise short-term volatility. From this, a statistical "**implied safety margin**" is derived, which represents the typical gap between the actual power flows and the boundary's operating (thermal) limit.




A trial smoothing algorithm (Exponentially Weighted Moving Average – EWMA) is then applied to the historical flow data. A revised implied safety margin is calculated for the smoothed flow. Three smoothing deployment cases are assessed to estimate the extent to which the reduced implied safety margin could allow additional energy to flow across the boundary, while maintaining the risk of exceeding boundary limits at historical levels.

Nested-boundary logic is used **to unwind BM bid-offer acceptances (BOAs)** where capacity exists to allow increased flows out of the constrained region. Potential cost and carbon savings are estimated using BM data and replacement generation carbon intensities (the amount of carbon dioxide produced by each type of generator). Finally, the analysis assesses the power and energy required of an asset providing smoothing in practice.

Key findings

- **Smoothing significantly reduces short-term fluctuations in power flow, which can lower volatility across the network.** However, the resulting increase in transferable energy is often limited as constraints elsewhere in the network prevent these improvements from translating into higher boundary flows.
- **The nested nature of boundaries can limit benefits,** as downstream constraints restrict the extent to which local reductions in safety margins deliver additional transferable energy.
- **Indicative gross BOA savings are very limited.** *Before accounting for service delivery costs,* observed annualised reductions across analysed boundaries are only in the range £0.7m to £4.0m.
- **Operational requirements are substantial.** Over the 15-month analysis period, the smoothing algorithm requires:
 - ~ 711 GWh of cumulative energy throughput.
 - ~ 3.3 GWh of effective storage capacity.
 - Frequent operation during constrained periods.
 - ~ 2.9 GW peak bi-directional power capability at individual boundaries.
- **Derived ceiling prices for a smoothing service are low relative to provider opportunity costs.** The values are estimated from avoided BOA costs and anticipated smoothing asset delivery requirements. They exclude the costs of service design, implementation and market operating costs. Estimated ceiling prices are:
 - £136 per hour of asset smoothing activity (£10 – £300).
 - £4.80 per MWh of energy moved to provide smoothing (£0.30 – £12.50).
 - £275 per MW of peak asset power capacity (£27 – £910).
 - £937 per MW of 10-minute-mean asset power capacity (£61 – £2440).



- £220 per MWh of asset energy storage capacity (£46 – £2940).

(Figures shown are weighted means with P10–P90 ranges in brackets)

- **These ceiling prices compare poorly with indicative historical battery energy storage system (BESS) revenues**, roughly **£57,000/MW-year** (power) and **£42,000/MWh-year** (energy), suggesting there would be very limited provider appetite under the evaluated setup.
- **Total carbon savings from smoothing are modest**, with estimated reductions of 1,220 to 6,080 tonnes per year across deployment scenarios, based on historical BOA data.

Assumptions and Limitations

This analysis provides an initial feasibility assessment and therefore adopts several idealised assumptions, many of which are likely to **overstate the potential benefits of smoothing**:


- **The implied safety margin may be overstated**, as its derivation assumes all variability is random, which increases the apparent scope for reduction by smoothing.
- **Smoothing asset performance is idealised**, assumed to provide instantaneous response and have both 100% efficiency and availability.
- **Network losses are omitted**.
- **Constraint BOA tagging is known to be incomplete**, which may result in either under- or over-estimate of impacts.
- **Independent treatment of boundaries overstates requirements**, as applying smoothing to each boundary separately over-estimates the volumes of energy and power flows required to deliver a smoothed flow.

Collectively, these assumptions are likely to overstate both the achievable savings and operational requirements. This could be refined in future work.

Technology Suitability and Opportunity Cost

Technologies with **fast response and high ramp rates** are best suited to delivering smoothing. **BESS, long-duration BESS and flexible data centres** meet these criteria. However, smoothing is highly intensive in terms of cycling, throughput and availability, which may limit a provider's ability to stack the service alongside other services.

As a result, the commercial proposition appears weak, with implied service ceiling prices an order of magnitude below typical historical asset earnings. This gap would widen further once full implementation, participation and delivery costs are taken into account.



Conclusions

Under current conditions and assumptions, **the economic case for boundary flow smoothing appears weak.** Gross BOA savings are limited, benefits are highly assumption sensitive, and nested constraints frequently restrict opportunities for realisable additional flow.

For service providers, the economics are also unfavourable: the inferred ceiling prices are low relative to both operational requirements and opportunity costs of alternative revenue streams.

From a bill payer perspective, once enabling and operational costs are included, a net benefit is unlikely.

2.Introduction


Project Context

Overview of Proposed Solution

Boundaries of Interest

Report Structure





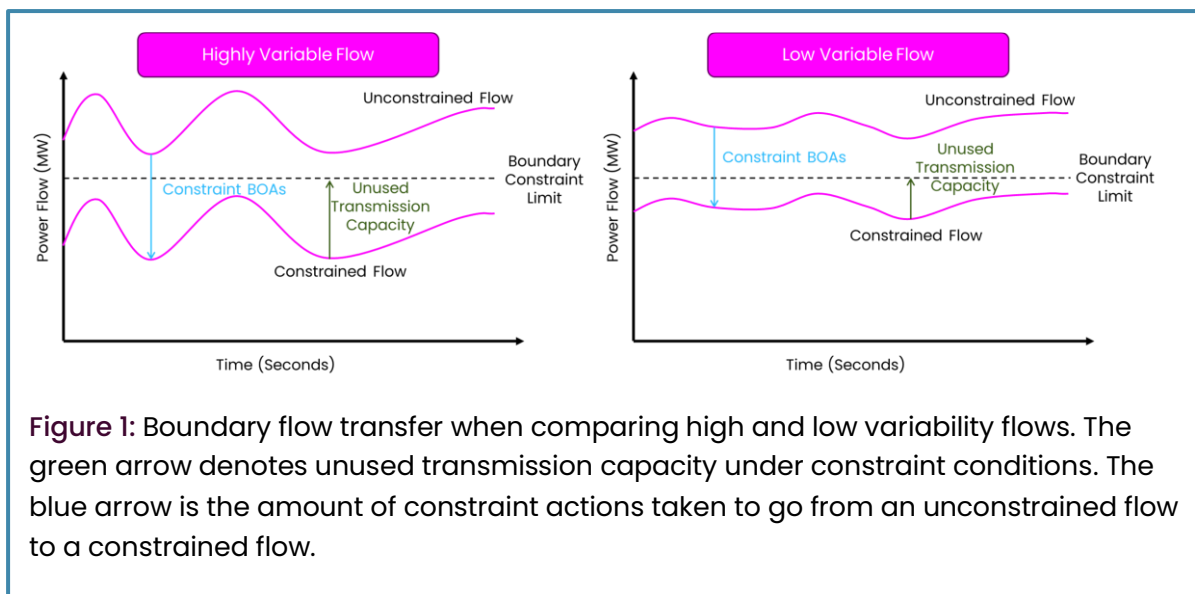

Problem Context

Power transfers across transmission boundaries in the Great Britain (GB) electricity network are governed by stringent operational constraints, primarily dictated by thermal ratings, system stability, and voltage limits. When these limits are approached or exceeded, system operators must implement remedial actions to maintain safe and secure system operation. These interventions are collectively referred to as constraint management.

Current constraint management practice relies predominantly on the constraint of low-carbon renewable generation that is located at the periphery of the grid, which is then often substituted with fossil-fuelled generation situated closer to demand centres. This approach incurs substantial operational cost. For example, in 2024, thermal constraint-related actions resulted in expenditures of £1.4 billion [1]. Such actions also impede progress toward national decarbonisation objectives. Transmission reinforcement represents the long-term solution; however, the planning, approval, and construction of new transmission infrastructure typically require several years, which limits its suitability for addressing short-term operational challenges.

The inherent volatility of boundary power flows introduces additional complexity. Rapid, large fluctuations in flow can result in frequent, short-timescale limit violations. Consequently, under conditions of elevated volatility, it may be necessary to reduce the permissible mean flow along a transmission corridor to ensure that peak excursions remain within operational limits.

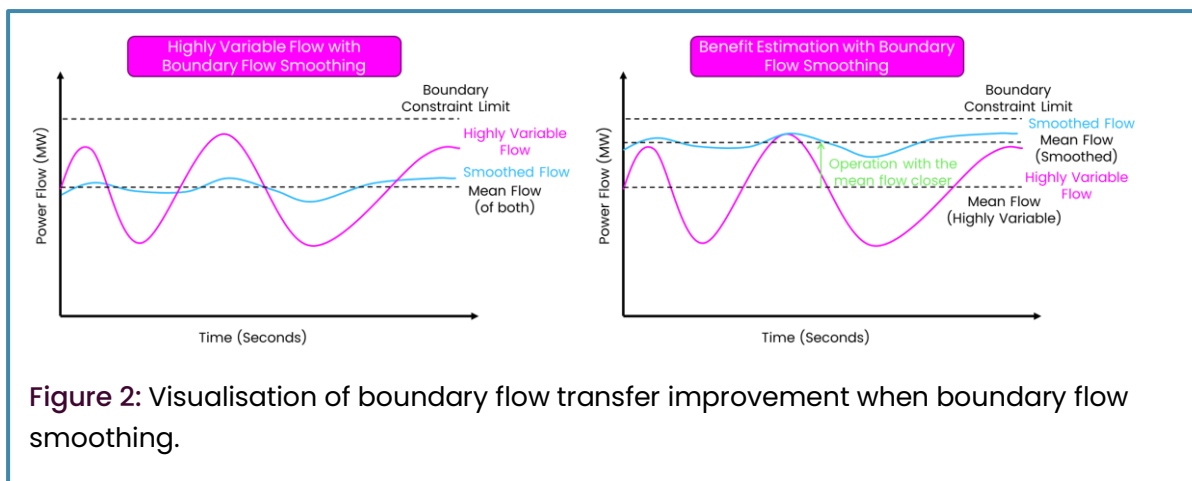

This study investigates the potential benefits of deploying assets designed to reduce short-term volatility in boundary flows, thereby enabling fewer constraint actions to be taken by the control room. The conceptual framework is illustrated in Figure 1, where the green arrow denotes unused transmission capacity under constraint conditions and the blue arrow is the amount of constraint actions taken (via BOAs) to go from an unconstrained flow to a constrained flow. In scenarios characterised by reduced volatility, this unused margin diminishes, allowing greater energy transfer across the boundary without breaching operational limits.



Actions to manage flows along transmission lines are linked to the wider processes of energy system management. System operation planning combines supply and demand forecasts, weather forecasts, outage plans and network modelling to produce a series of assessments that are updated closer to real time as further information is received. The Control Room (Electricity National Control Centre (ENCC) who are responsible for operating the grid) receives updated outputs from a Power Network Analysis (PNA) every 10-minutes and has sight of live power flows and voltages across the transmission network. These are used to identify violations of voltage or thermal limits on the network, which then inform the adjustment of constraint limits. The updated limits inform necessary Balancing Mechanism actions to keep flows within prescribed limits in close to real time. The planning process is discussed in greater detail in 'Appendix A – Current Operational Practices'.

Overview of Proposed Solution

A boundary flow smoothing service could work as shown in Figure 2. Here smoothing has reduced the volatility in power flows. The less peaky flow means that fewer constraint actions are required, and the flow can be operated closer to the limit, reducing the costs of constraint management. The green arrow represents the increased transmission capacity usage from the reduction in volatility from boundary flow smoothing. For boundary flow smoothing to be a viable operational service, these cost savings must exceed the expenses associated with implementing the scheme, operating the control mechanism, and contracting the assets that deliver the smoothing capability.



Boundaries of Interest


This work has focussed on demonstrating the feasibility of boundary flow smoothing on the most constrained and costly-to-manage sections of the GB transmission system, specifically the set of boundaries spanning Scotland and northern England. These boundaries provide a suitable test case because the associated constraint costs make potential savings easier to demonstrate. However, the methodology developed here is not specific to this region. The same approach can be applied to any transmission boundary where short-term flow volatility contributes to operational constraints and associated costs.

These are set out in the box below and their locations shown in Figure 3.

The following boundaries are considered within this project:

1. SSEN-S (B2) – North-to-South SSEN Transmission
2. SSE-SP2 (B4) – SSEN Transmission to SP Transmission boundary (shared by SSEN Transmission and SP Transmission)
3. NKILGRMO (B5) – North to South SP Transmission
4. SCOTEX (B6) – SP Transmission to NGET (shared by SP Transmission and National Grid Electricity Transmission)
5. HARSPNBL (B6a)
6. SSHARN-3 (B7a) – Upper North of England

The boundaries of interest form a set of so-called nested boundaries. Because the power transfer is predominantly north to south crossing all boundaries, actions taken on a



northern boundary propagate southwards, to produce effects on all downstream boundaries.

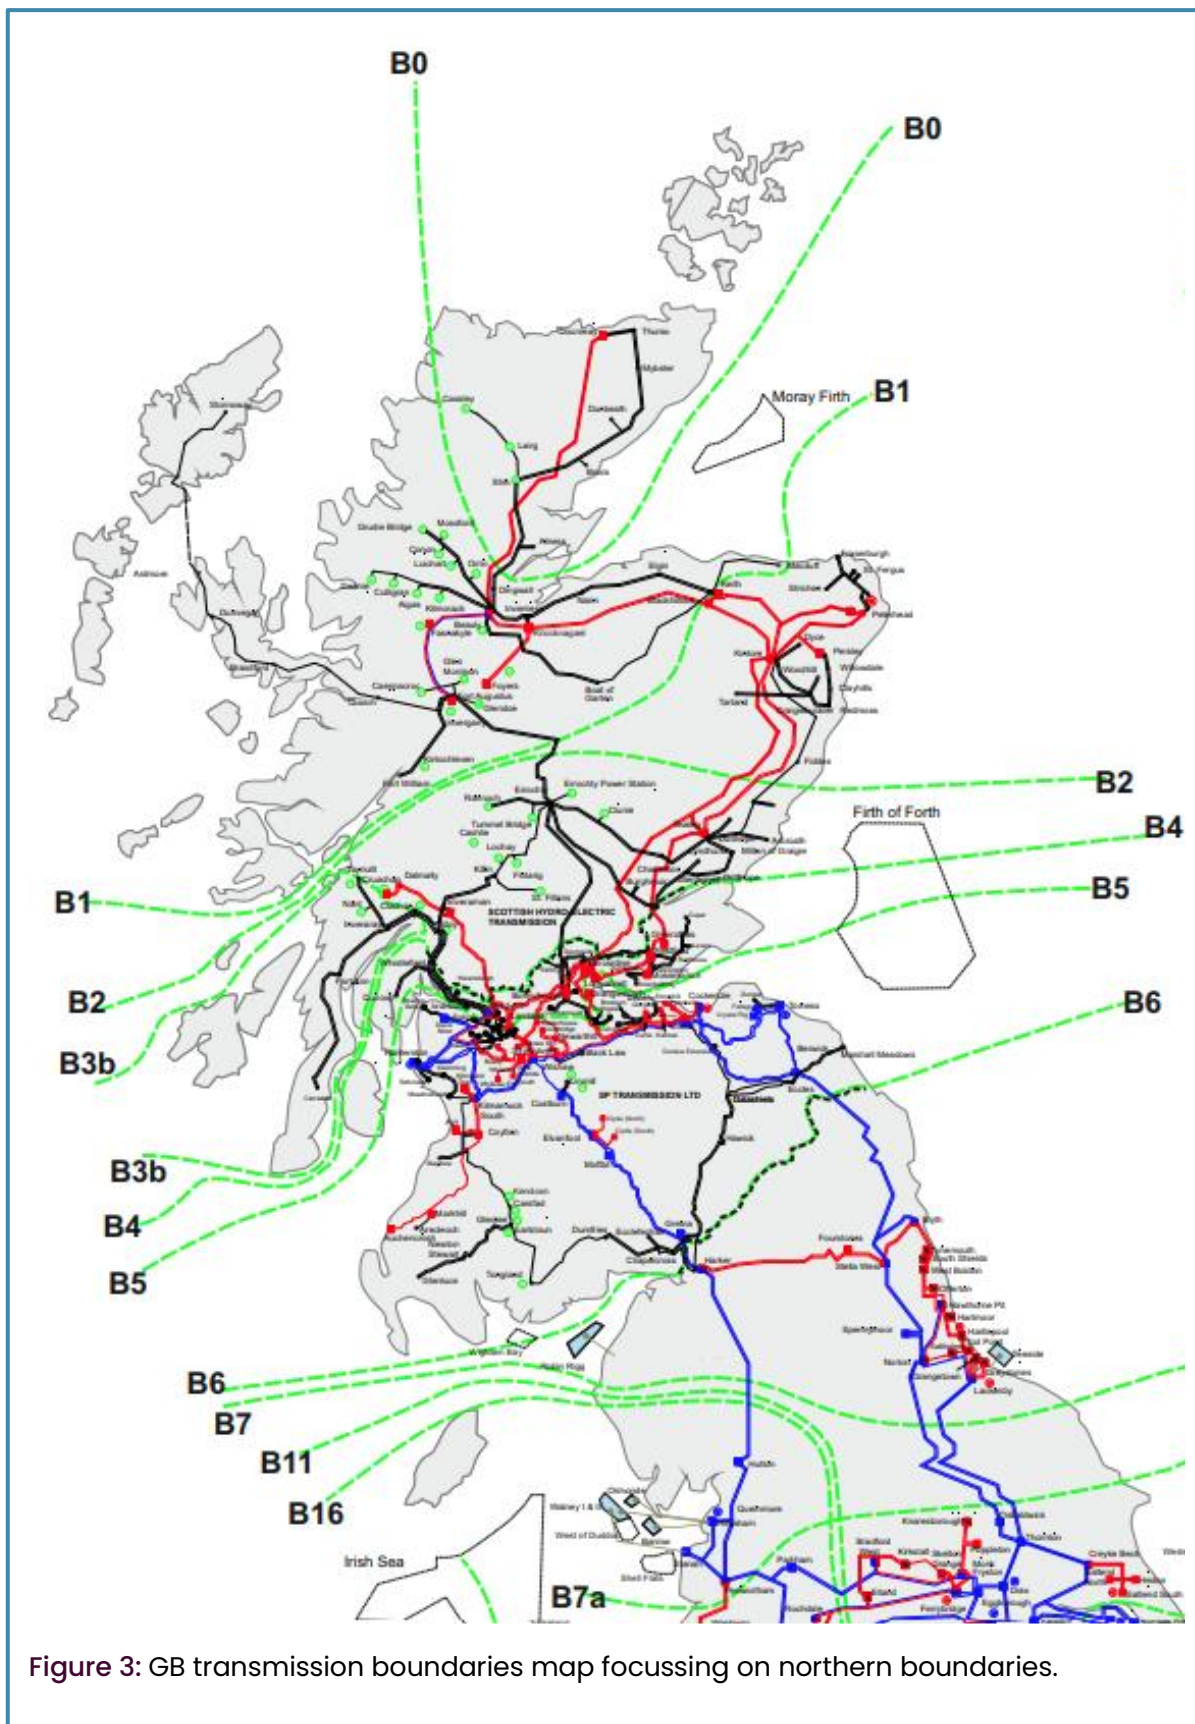



Figure 3: GB transmission boundaries map focussing on northern boundaries.



Report Structure

This report covers the following areas:

- Section 3: Research and Theoretical Feasibility
 - What existing examples have been used to enable higher boundary flows?
 - What does the historical boundary flow behaviour look like?
 - How can boundary flows be statistically quantified?
- Section 4: Cost-Benefit Analysis Approach
 - How are smoothing benefits realised?
 - How can the costs and benefits be quantified?
 - How should the Cost-Benefit Analysis (CBA) be structured and what should be included?
- Section 5: Cost-Benefit Analysis Results
 - What are the results of the CBA?
- Section 6: Discussion and Conclusions
 - What do the results mean?
 - What assumptions have been used?
 - What technologies are suitable for smoothing?
 - What could be explored further?

The CBA results and discussion aim to answer:


1. Can smoothing boundary flows enable higher transfer capacity without breaching limits?
2. What algorithms and technologies can deliver effective smoothing?
3. How many constraint actions could be avoided, and what are the associated cost savings?

3. Research and Theoretical Feasibility

International Comparison Projects

Historical Boundary Flow Behaviour





Section Overview

This section presents a review of schemes used in other countries and an analysis of historical boundary flows to gauge the theoretical feasibility of the concept and potential for capacity savings.

International Comparison Projects


A review of academic literature and industry publications revealed no documented precedents for this specific form of transmission line flow smoothing. Although no directly comparable interventions were identified, several international projects demonstrate the use of energy storage to support transmission networks. These initiatives typically rely on the virtual transmission line (VTL) concept, in which storage assets operate as transmission-supporting equipment that can enhance transfer capability, provide redundancy, and offer short-term relief from constraints. This is known as using storage as transmission assets.

The literature review therefore indicates that, while there are numerous international examples of storage being deployed to reinforce transmission capacity, there are no published cases of storage being used explicitly to smooth transmission line flows in order to reduce constraint costs. Because the objectives of existing VTL projects differ substantially from those of boundary flow smoothing, they provide limited transferable lessons for this study.

The literature review has highlighted the following:

- VTL systems have been successfully implemented in several transmission networks worldwide, with additional schemes planned.
- Many of the existing projects operate primarily during fault or outage conditions, whereas the boundary flow smoothing concept is intended for routine operational use.
- Storage as transmission assets could be deployed in GB to boost transmission line capacity. Existing examples have assets owned and operated by Transmission System Operators (TSOs), the ownership and operation of storage as transmission assets would have to be considered for a GB context.
- The timeline for building storage as transmission assets may not be quick enough to bring benefits before planned transmission capacity upgrades. The Constraints Collaboration Project (CCP) is seeking short-medium term solutions that would deliver significant consumer value ahead of network reinforcement, therefore the focus within this project is based on the use of existing rather than planned flexible assets.

The development timescales for storage as transmission assets may be comparable to those of traditional network reinforcement, which limits their suitability for delivering



near-term constraint mitigation. By contrast, boundary flow smoothing is conceived as a short-term operational measure.

Overall, the reviewed projects demonstrate the feasibility of using storage to support transmission networks, but they differ fundamentally from the proposed smoothing approach, which seeks to mitigate routine operational fluctuations rather than respond to outages or faults. Further details of the literature review can be found in 'Appendix B – Literature Review'.

Historical Boundary Flow Behaviour

As part of this project, second-by-second boundary power flow and constraint data for the nested set described in the **Boundaries of Interest** section was analysed for the period 01/01/2024 to 31/03/2025 (to cover 2 winter periods, as these are typically when constraints are the worst). Table 1 provides a summary of the metered flow data during constrained periods, which highlights the large amounts of energy turned down (bids) and the large associated cost to these actions. The following terms have been defined to describe the flow and have been applied throughout this report:


- **Adjusted flow** is the power flow with the bids and offers unwound from the metered flow to calculate a flow that precedes any actions taken by the control room e.g. constraint actions.
- **Constrained periods** are defined as when the adjusted flow was above the constraint limit.
- **Volatility** refers to short-term stochastic fluctuations in boundary power flow.
- **Exceedance** refers to how much energy is transferred above the constraint limit i.e. how much and how often the constraint limit is exceeded.

A more in-depth overview of the data and this initial analysis is provided in 'Appendix C – Data Overview' and 'Appendix D – Understanding the Scale of Boundary Flows and Constraints'.

Boundary flow volatility is influenced by several predictable macro-level drivers that affect both supply and demand. These include:

- Seasonality
- Weather conditions
- Daily demand profiles

Superimposed on these drivers are short-timescale fluctuations, which could be considered as noise. These originate primarily from variations in wind generation output and rapid changes in demand. Reducing this component of boundary power flow volatility would simplify the management of constraint limits by increasing control room



confidence that flows will remain within the transfer limits established to ensure system safety and stability.

For the statistical analysis of boundary flows within this project, the data were separated into two components: the volatility within 10-minute windows; and the variation that exists between the 10-minute mean flows of these windows. Splitting the flow into 10-minute windows gives 65,520 instances of 10-minute means and variance data. The results of the analysis of this data are presented in the **10-Minute Averages – Δ (P90)** section and **10-Minute Averages – Evolution of the Mean** section. Further details for each boundary may be found in ‘Appendix E – Distribution of Mean Flow’.

This approach is justified for three reasons. First, the control room updates boundary constraints on a 10-minute cycle, so the chosen resolution aligns the analysis with operational practice. Second, a 10-minute interval corresponds to the typical response time for operational actions such as generator notifications and ramping, which ensures that the analysis reflects the timescales over which corrective measures take effect. Third, flow limits are generally governed by thermal considerations, and transmission assets have 10-minute thermal ratings (they also have additional thermal ratings for other time periods) that specify the duration for which they can safely withstand an overload before the onset of thermal damage.

A more comprehensive exploration of boundary flow volatility was conducted using a wider array of statistical approaches. These included:

- Range and variance
- Spectral techniques (Fourier analysis)
- Realised volatility

After investigation these have not been used to assess the smoothing effect and instead the statistical approach outlined above was chosen as a better approach. However, some of these did feed into the smoothing assessment metrics that were used. A description of the analysis and results is detailed in ‘Appendix F – Variability and Volatility Analysis’.





Table 1: Metered flow data for constrained periods during analysis period (01/01/24 to 31/03/25)

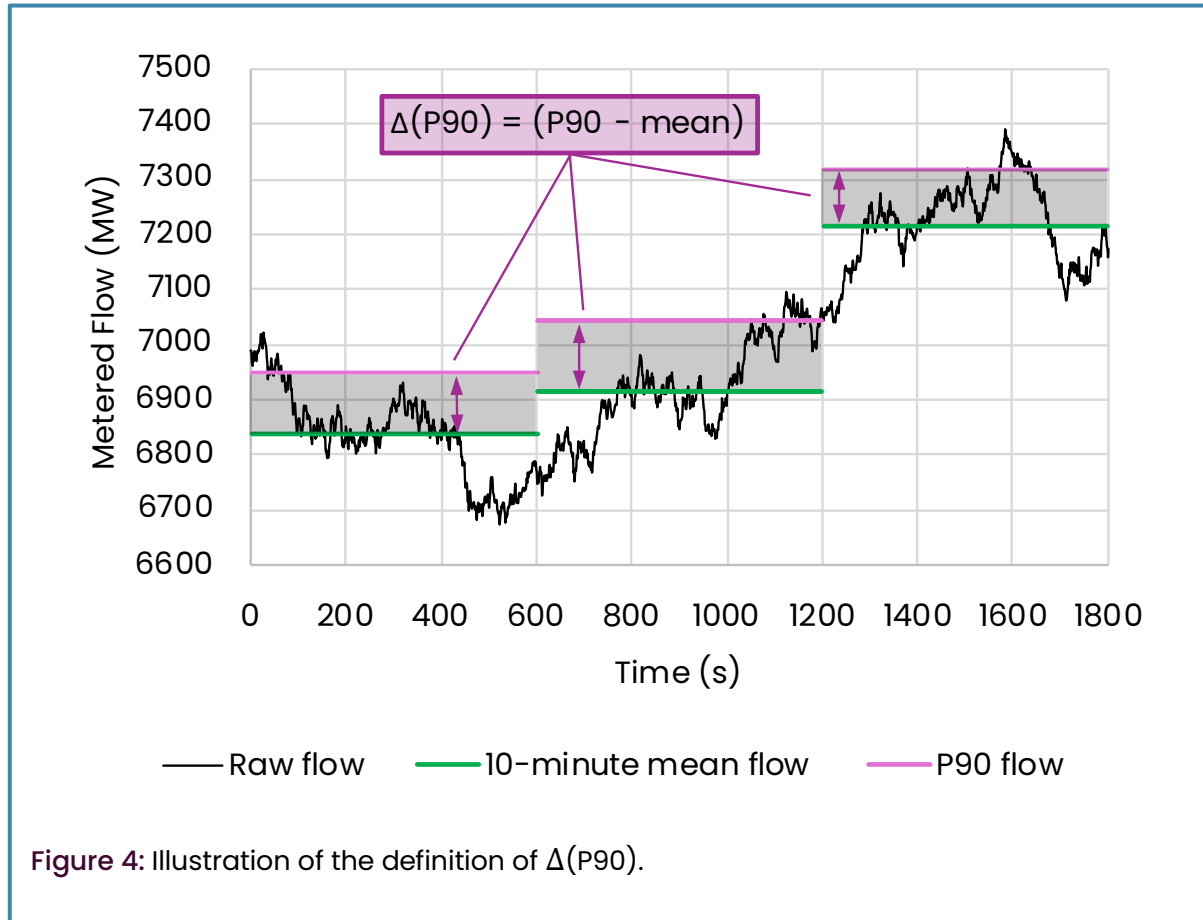
Boundary	Metered Flow (MWh)	Limit (MWh)	Difference (Flow – Limit) (MWh)	Energy Transferred in Exceedance of the Flow Limit (MWh)	Action Cost (£)	Constrained Energy¹ (MWh)	Time Above Limit (days)	Percentage of Time Above Limit for Constrained Periods
SSEN-S	4,680,000	5,130,000	-450,000	12,600	83,100,000	2,400,000	10.90	10.80%
SSE-SP2	7,530,000	8,100,000	-570,000	73,600	68,300,000	2,830,000	40.00	24.80%
NKILGRMO	3,740,000	4,650,000	-910,000	1,040	5,320,000	686,000	0.81	0.01%
SCOTEX	8,290,000	8,640,000	-350,000	411,000	26,600,000	562,000	16.60	21.50%
HARSPNBLY	11,100,000	11,400,000	-300,000	309,000	16,600,000	467,000	26.70	20.90%
SSHARN3	12,900,000	13,500,000	-600,000	18,400	19,300,000	344,000	6.90	9.90%

¹ Calculated by summing the total energy of constraint-tagged BOAs

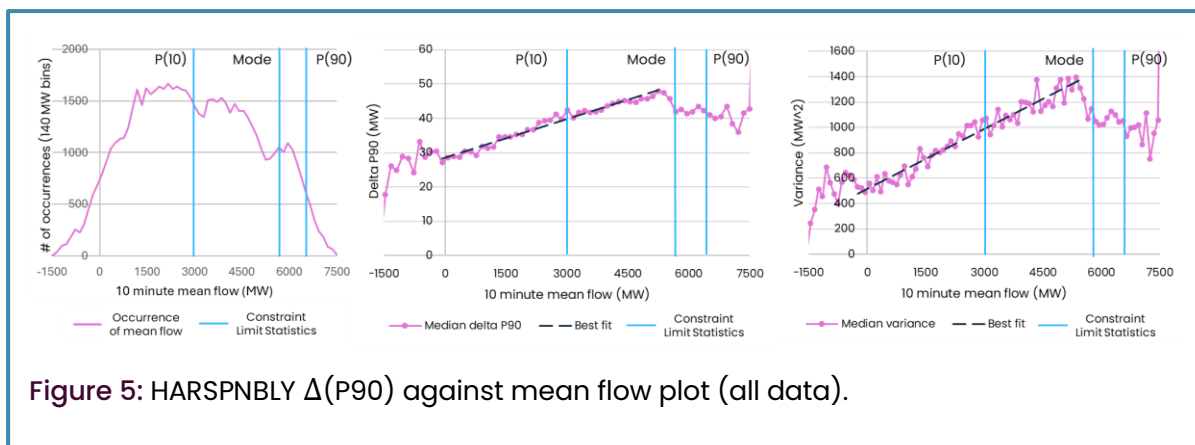



10-Minute Averages – $\Delta(P90)$

The volatility within each 10-minute window is characterised using two statistical measures: the standard deviation and the difference between the 90th percentile flow and the mean, which is denoted as $\Delta(P90)$. This is illustrated in Figure 4. The 90th percentile flow (P90) is defined as the flow value that exceeds 90% of all observations within the 10-minute interval. In other words, the flow is above the P90 threshold for a cumulative duration of approximately one minute, although these exceedances need not be contiguous measurements.



For each of the ~65,000 10-minute windows, a value for the mean flow, flow variance and $\Delta(P90)$ is calculated. By binning the data according to the mean flow value, it is possible to identify how the volatility of flow changes with flow magnitude. The outcome of this is presented for the HARSPNBLY boundary in Figure 5. This shows the occurrence of different flow magnitudes, the median value of $\Delta(P90)$ (MW) within each bin and the variance of the flow, each plotted as a function of the mean flow (MW) across the HARSPNBLY boundary. The (P10), mode and (P90) of the flow limits is plotted as the blue lines. Comparable patterns were identified across all other boundaries examined in this study.



The occurrence of mean flow indicates that the most typical flows occur between 1000 and 5000 MW. There is a large plateau peak that occurs between 500 and 6500 MW with the occurrences dropping off rapidly to a minimum of -1500 MW and a maximum of 7500 MW.

At low and moderate power levels, both median $\Delta(P90)$ and the variance increase approximately linearly with mean power flow. This relationship weakens and subsequently plateaus or falls at higher flow levels, with the flow $\Delta(P90)$ and variance peaking typically about 5000 MW, before the modal flow limit value. Similar patterns were observed across all boundaries included in the analysis.


The increase in volatility as flow increases is not unexpected. For these boundaries in particular, higher flows are assumed to indicate higher output from the large amount of wind generation located offshore in Scotland. On a sub-10-minute timescale, fluctuations in windfarm outputs will not be correlated, meaning that the variances of their individual flows can be summed to give the variance of the total flow.

There could be a number of possible reasons for the trend seen in the variability peaking and then plateauing or subsiding at higher flows. Firstly, this is observed to occur as the flow limit is approached. This could indicate that wind farms are being constrained and therefore the observed flows are coming from fewer windfarms, meaning that the volatility would fall. Secondly, high levels of wind generation could need other generation to come online that is less volatile for grid stability purposes. Thirdly, windfarm output curves show sharp increases in output at low-moderate wind speeds and then plateau so, at higher wind speeds, the output from an individual wind turbine will be constant and less volatile.

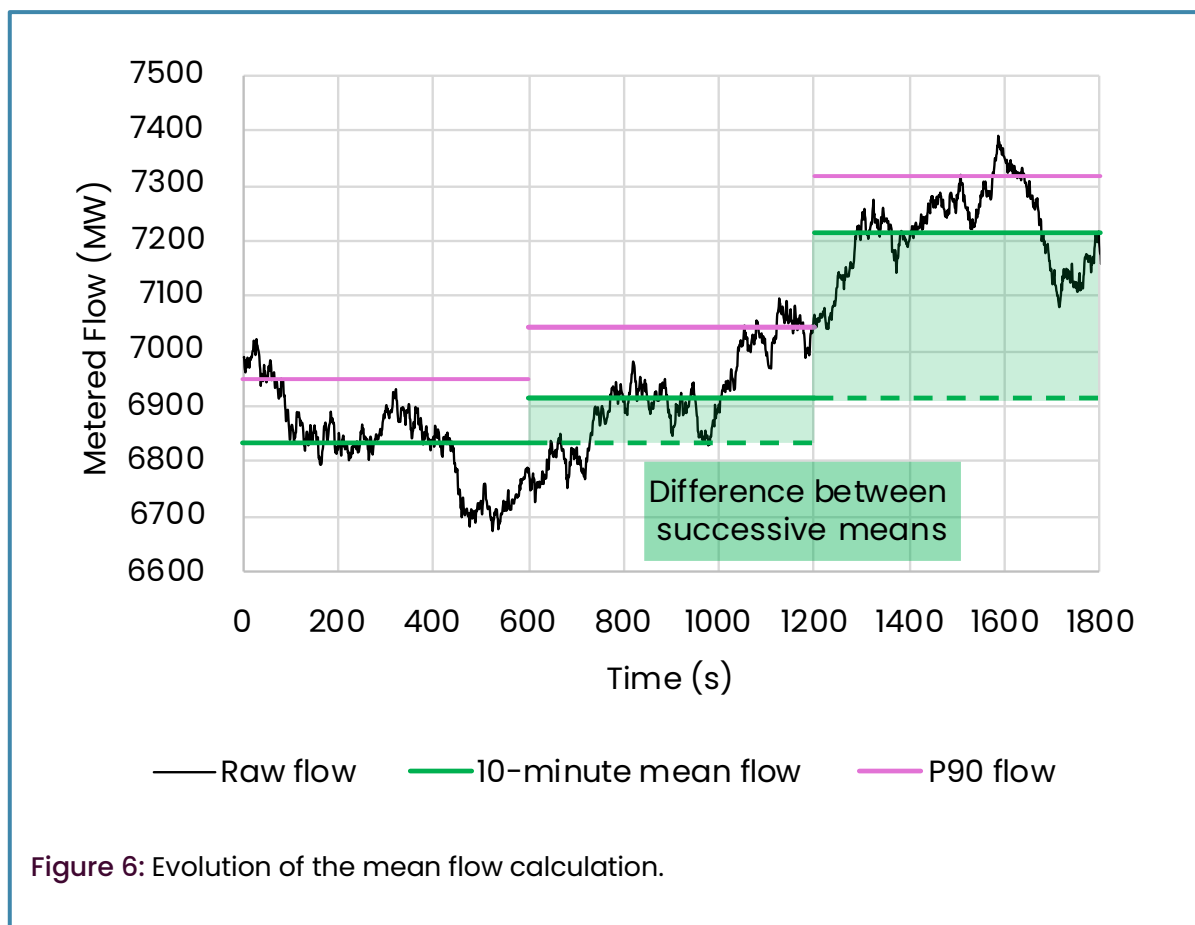
This project is concerned with the power flows in the vicinity of the flow limit, as outside of this smoothing provides no additional benefit on top of existing operational mechanisms such as the BM. The above analysis indicates that the volatility of the flows of interest (high levels) show little correlation with flow levels.

10-Minute Averages – Evolution of the Mean

The previous section highlights the variation of the boundary flows relative to the mean within 10-minute windows. This section describes how the 10-minute mean values evolve.



Insight into this is gained by looking at the statistics of the differences between adjacent successive values – a quantity known as the first difference, illustrated in Figure 6.




First difference data is calculated and binned according to its size to derive a statistical distribution that describes the likelihood that the mean in the next 10 minutes will differ from the mean in the previous 10 minutes by a certain amount. Distributions can be created for all 10-minute periods, or periods can be filtered according to characteristics of the flow. For example, periods where the flow was high, periods that were constrained or periods that were unconstrained.

Figure 7 shows a number of distributions for the flow data sets at SSE-SP2.

- “All Periods” includes all the metered power flow data.
- “Consecutive Periods Not Constrained” includes metered data where successive 10-minute periods were both not constrained.
- “Consecutive Periods Constrained” includes metered data where successive 10-minute periods were both constrained.
- “Consecutive Periods Different” includes metered data where successive 10-minute periods were different i.e. one was constrained and the other was not.

The distributions for all the data sets differ but in particular “Consecutive Periods Different” is noticeably different which suggests the behaviour changes as the flow moves from constrained to not constrained and vice versa. This is to be expected as constrained



periods beginning or ending would indicate generators being constrained or unconstrained so step changes in flow would be seen rather than typical power flow noise.

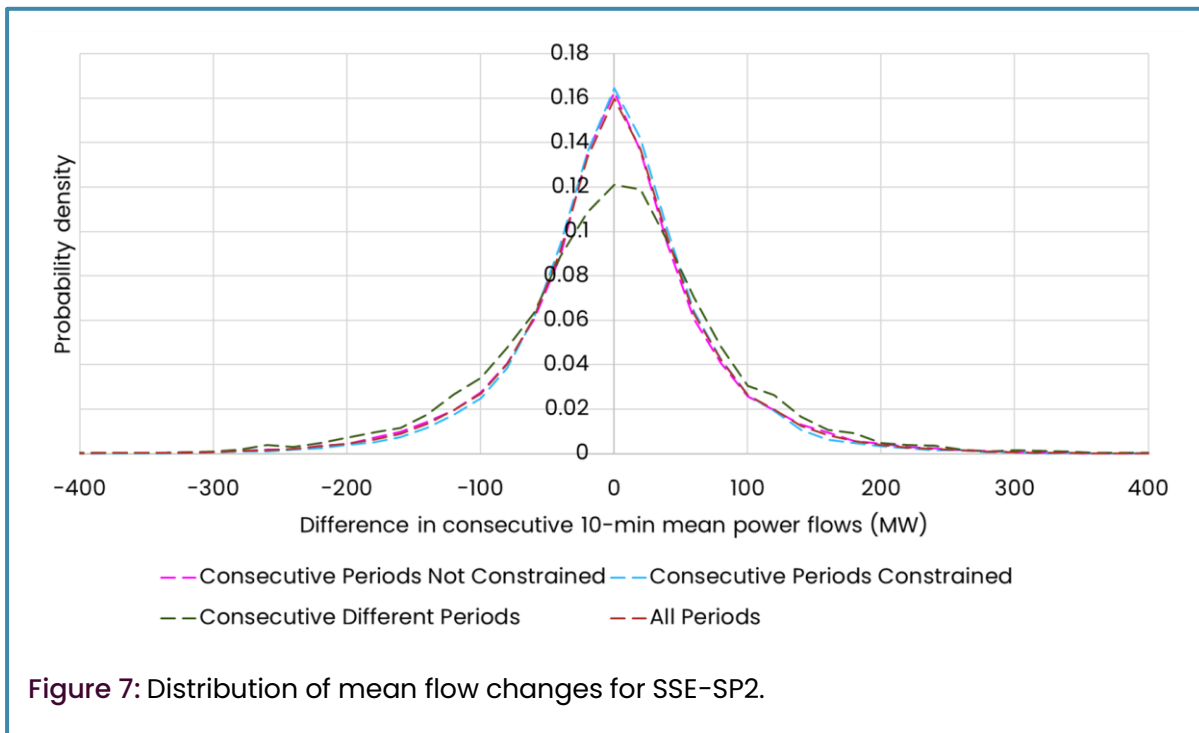
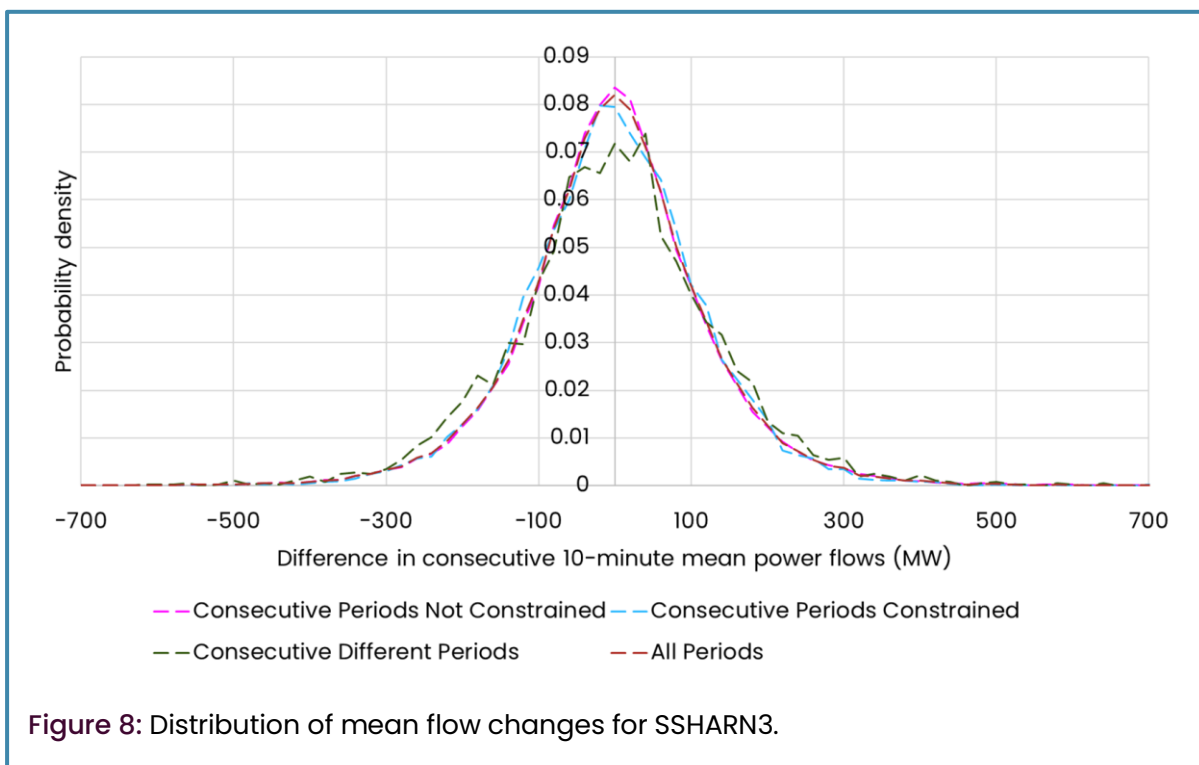



Figure 8 shows a number of distributions for the flow data sets at SSHARN-3. The analysis shows similar conclusions for SSE-SP2, in that the distributions differ in particular “Consecutive Different Periods”.





This analysis shows that, at least over time periods shorter than those typically influenced by macro factors, the evolution of historical flows across boundaries changes in a way that can be modelled by definable statistical distributions. The characteristics of these distributions change depending on whether periods are constrained or not. With the distributions for flow data, statistics can be produced on the likelihood that smoothed or unsmoothed data will exceed a limit within a certain timeframe.

A wider analysis of various power flow data sets was undertaken as part of the project. The difference in the behaviour of power flows was explored for constrained periods, unconstrained and transitional periods. It was found that there is essentially zero probability that the three subset distributions are from the same underlying distribution².

These results indicate that:

- Control room constraint actions do change the variance statistics of the flow when it is considered in 10-minute blocks. Therefore, when considering the impact of smoothing algorithms on flows, the performance for constrained and unconstrained periods should be separately assessed.
- Historical boundary flow data can be used to gain insight into the likely statistical behaviour of boundary flows. The analysis undertaken in the project has shown that, statistically, flows behave differently depending on whether control room constraint actions were taken. At high power flows, there are sufficient instances of both constrained and unconstrained periods to allow us to determine for both cases the statistics of how the mean flow evolves and how flow fluctuates upwards from this. These two factors will allow the prediction of how likely smoothing is to deliver additional, usable capacity on constrained lines.

² The adjusted flows (all BOAs unwound from the metered flow see 'Appendix C – Data Overview') were also compared and showed markedly different behaviour to the metered flow, both for constrained and unconstrained periods. There is essentially zero probability that these are drawn from the same distribution. Much heavier tails indicate that, through removing the bids and offers to generate the adjusted flows, significantly larger swings in power between successive 10-minute periods are generated. It was decided that the adjusted flows may not be a suitable 'counterfactual' flow and should not be used to test smoothing algorithms, as they do not give a realistic representation of boundary flow volatility.

4. Cost Benefit Analysis

Approach


Introduction

Routes to Smoothing Benefits

Quantification of Smoothing Costs and Benefits

Cost Benefit Analysis





Introduction

The purpose of the CBA is to estimate the level of savings that might be achieved through the use of assets to smooth cross boundary flows, and to establish whether this represents an economically feasible way to reduce constraint costs in the short term.

The CBA compares two cases:

- A business-as-usual (counterfactual) case, based on historical boundary flows between 01/01/24 and 31/03/25.
- A smoothed flow case, in which historical flows are revised based on the statistical effects of smoothing and on unwinding historical BOA decisions.

The benefits of smoothing are captured by estimating the volume of BOAs in the business-as-usual case that can be unwound in the smoothed flow case. This allows a calculation of the monetary savings made by smoothing. The BM savings calculated in this report are based on actual historical bid and offer prices paid during the period studied – i.e. Jan 2024 to March 2025.


The full costs of smoothing include:

- Implementing communications.
- Developing robust, asset-specific algorithms that do not conflict with other system-management processes (e.g. dynamic frequency response services).
- Establishing and running a market.
- Paying service providers for delivery or availability.

At this early feasibility study stage, these costs are unknown. To manage this, the analysis derives ceiling prices for the service (per hour, per MW and per MWh) through combining BOA cost savings with the power and energy flows required for smoothing. The implications of the power and energy flows for asset specification are calculated, which allows an estimate of the annual asset earnings linked to asset capacities.

The CBA approach covers the following:

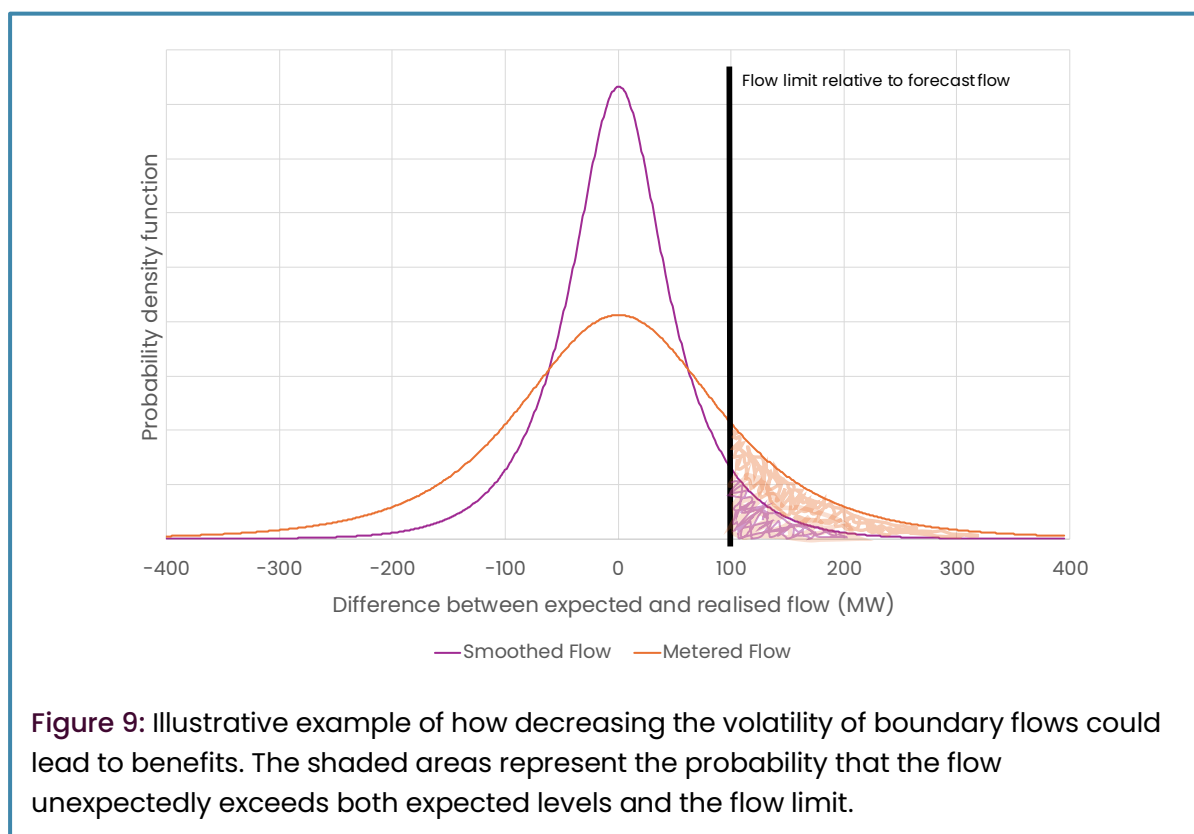
- Routes to smoothing benefits – what benefits can be realised from smoothing?
- Method for quantification of smoothing benefits and energy/power flow requirements:
 - Selection of a trial smoothing algorithm.
 - Derivation of safety margins.
 - Application of flow increases across nested boundaries.
 - Quantification of the cost savings due to reduced BOA actions.



- Calculation of the requirements from smoothing assets.
- Approaches to combining costs and benefits.


Routes to Smoothing Benefits

Figure 9 below compares the distributions of discrepancies between the expected and realised boundary flow for illustrative smoothed and unsmoothed cases. The probability of large deviations from expected or forecast values is lower under the smoothed profile, indicated by the smaller size of the pink-shaded area as compared with the orange-shaded area. This section will discuss qualitatively the benefits that could arise from this increased certainty.



Improved Asset Reliability

Assuming the mean flow remains the same, a less volatile flow brings direct benefits through reducing peak power flow spikes and boundary flow limit breaches. Boundary flow limits will often be calculated from equipment thermal limits, so exceedance of these limits has the potential to cause overrating damage. However, the transmission system is operated on the “N-1-1” requirement where the system is secured against faults, so it is unlikely that thermal limits will be exceeded in non-fault scenarios. The decreased volatility from smoothing leads to potential increases in asset lifetimes and reduced risk of failure. Although these are the direct benefits of a smoothed flow, they are not the



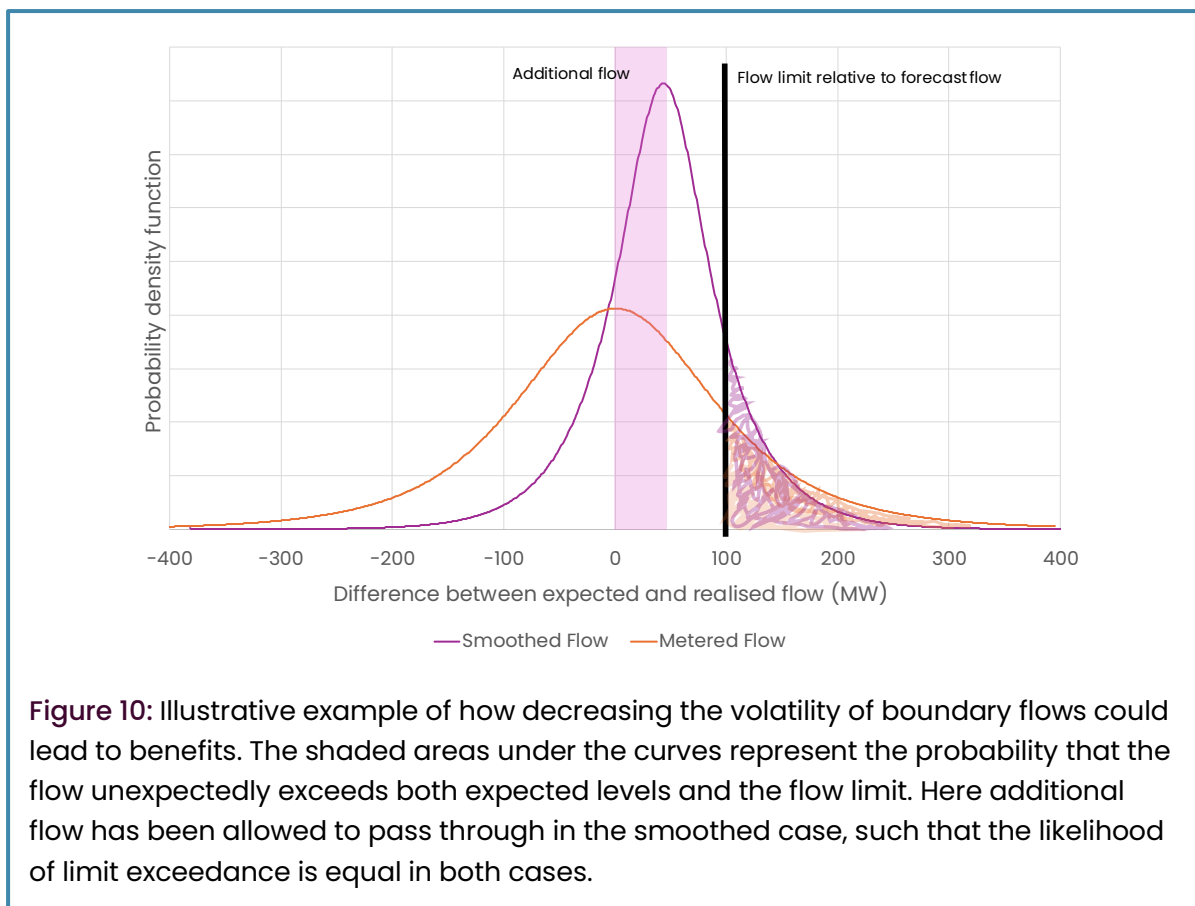

motivation for this proposed intervention, and their potential value will not be calculated here. The analysis will aim to keep the level of constraint limit exceedances the same for the smoothed case compared to historical levels. This ensures that this uncalculated, additional benefit has minimal impact on results.

Reduced Requirement for Constraint Actions

The primary mechanism used by NESO to manage transmission constraints is the BM. Within the BM, NESO may instruct balancing mechanism units (BMUs) located behind a constraint to reduce their output by accepting their bids and instruct BMUs located ahead of the constraint to increase output by accepting their offers, thereby maintaining overall system balance. Each instruction issued through this process is referred to as a BOA.

Smoothing the power flow reduces short-term volatility and increases confidence that flows will remain within operational limits. However, it is important to note that smoothing does not materially change the underlying mean flow across the boundary³. Since the average level of power transfer is likely to remain the same, smoothing alone cannot reduce the overall requirement for constraint actions. However, the increased certainty that smoothing offers may enable the control room to adopt less conservative BOA strategies. In effect, some of the benefits of reduced peak flow and lower exceedance likelihood could be traded for the opportunity to transmit additional power across constrained boundaries. This is illustrated in Figure 10 below.

³ This is because energy stored in the smoothing asset is returned to the system on timescales that are short compared to the typical duration of exceedances. Flow across the smoothed boundary will change over the whole extent of the analysis by the difference in state of charge of the asset at time zero and the end. Also in reality, though not included in this analysis, the inefficiency of the smoothing asset would also reduce flows slightly.





For smoothing to deliver this benefit, there must be a mechanism that links Control Room decisions on constraint management to the volatility of the flow, i.e. to the level of confidence in future flow conditions. In other words, an appreciation of volatility on short timescales needs to be a part of the decision-making framework to assess how operational outcomes might differ if volatility were to be reduced.

The outcomes of current decision-making practices are evident from observed metered flow data and the associated BOA records. However, it is unclear whether, or in what way, a smoothed flow would alter these decisions, since there is no formal BOA selection process that explicitly incorporates flow volatility.

To address this gap, the following operational approach is assumed:

The number of constraint limit exceedances is controlled by leaving an implied safety margin that accommodates unexpected upward fluctuations in flow.

So, while NESO do not operate with an official safety margin, the analysis assumes that this is informally adopted into decisions. An 'acceptable' number of exceedances is established by observing the historical occurrence of exceedances during periods of constraint at each boundary. The concept of the implied safety margin is illustrated in Figure 11.



Smoothing the flow would narrow the distribution of realised flows in Figure 9, as per the illustration in Figure 10, and permit fewer BOAs to be taken while matching the probability of limit exceedance. Matching the probability of limit exceedance to historical levels ensures that the risks of damage due to exceedance is roughly equivalent in both cases.

The Impact of Nested Boundaries

The situation is made more complex by the nesting of boundaries. In reality, this means both that:

- Some of the benefit of smoothing at a boundary is transmitted southwards to other boundaries.
- Options for taking alternative BOA decisions as a result of smoothing may be restricted due to the location of units that are bid off.

The first effect is omitted from this analysis, though could be considered in further work.

The second effect is accounted for in the analysis so that a set of revised flows across all the boundaries is calculated that reflects any actions taken at each boundary. This allows the smoothing benefits to be calculated for the full set of boundaries described in Boundaries of Interest.

Method for Quantifying Smoothing Costs and Benefits

This section describes the approach used to quantify the reduction in BOA requirements. Figure 12 outlines the process. The purple circles indicate analysis or processing and are described in greater detail in the following subsections referenced by the bracketed numbers.

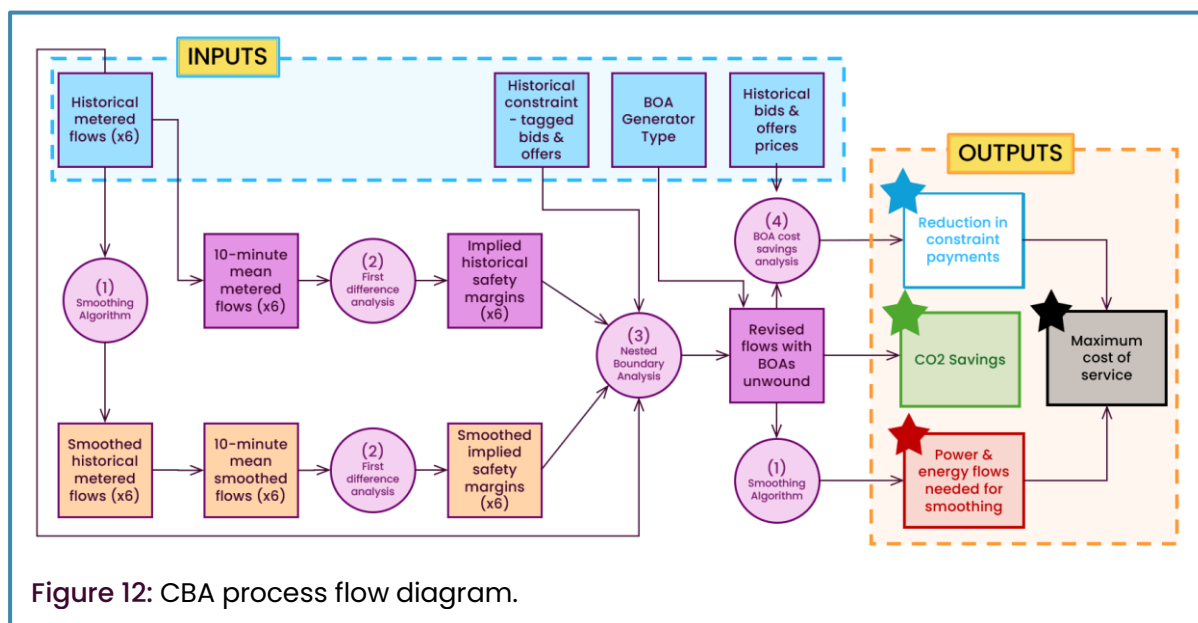



Figure 12: CBA process flow diagram.

In summary:

- A suitable trial algorithm for smoothing the flows is identified.
- Analysis of first differences (as described in [10-Minute Averages – Evolution of the Mean](#)) is applied to the historical flow data to infer the business-as-usual implied safety margin for each boundary.
- The historical flow data is smoothed using the chosen smoothing algorithm, and first difference analysis is applied to derive an implied safety margin for smoothed flow at each boundary.
- The historical flows, BOA data and the two implied safety margin values are combined in an analysis of the nested boundary system to provide a revised set of BOA actions in which:
 - Bids are unwound within the nested region where smoothing allows for fewer constraint actions to be taken.
- Corresponding replacement energy offers are unwound south of SSHARN3, the southern-most of the boundaries considered here.



- The volumes of BOAs unwound at each boundary are combined with BM data on the costs per MWh of actions taken, the type of generators involved and their carbon intensities allowing both total cost and carbon savings to be estimated. This represents the benefits of smoothing.
- The revised flows determined by the unwound BOAs are passed through the smoothing algorithm to determine the flows of power and energy that were required to deliver the smoothing. This provides an estimate of the break-even cost of providing a smoothing service.

The remainder of this section provides greater detail on these steps.

Identification of Trial Smoothing Algorithms (1)

An initial list of potential smoothing algorithms was identified and reviewed to understand expected behaviour and suitability. Qualitative assessment of how effectively each algorithm handled typical data patterns was undertaken to provide a shortlist.

The algorithms in this shortlist were then evaluated quantitatively. To undertake this, each algorithm was consistently applied to the boundary flow dataset covering 15 months of operational data. The resulting outputs were compared to determine how well each algorithm performed against a set of assessment metrics measuring smoothness, constraint limit avoidance and asset energy usage.

This analysis showed that the exponentially weighted moving average (EWMA) provided the strongest overall performance. It offered an effective balance between smoothing quality and asset usage, making it the most appropriate option.

Full details of the smoothing algorithm analysis is contained within 'Appendix G – Smoothing Algorithms' and the algorithm optimisation results are shown in 'Appendix H – Algorithm Optimisation Results'.

First Difference Analysis & Reduction in Implied Safety Margin (2)


Inputs: Historical flow data for 6 boundaries; Choice of smoothing algorithm.

Outputs: Reduction in implied safety margin required to maintain historical levels of limit exceedance in the smoothed flow.

This section describes the approach adopted for estimating the implied safety margin, based on the distributions of the first difference values, described in the [10-Minute Averages – Evolution of the Mean](#) section.

Basis for Approach

The onward analysis assumes that Control Room constraint decisions are based on unbiased forecasts (i.e., forecasts that are on average neither too high nor too low). The actual flow realised in any 10-minute period can be thought of as the forecast value plus a



noise component; here, noise refers to the irreducible, stochastic component of short-term flow variability rather than any form of forecast inaccuracy. To represent this noise, the first-difference distributions (from the [10-Minute Averages – Evolution of the Mean](#) section) are used. This is appropriate because:

- The empirical distribution has the characteristics expected of short-term stochastic noise in time-series analysis, being approximately bell-shaped and symmetric around a modal change of 0 MW (i.e. the most likely 10-minute change is 0 MW).
- Macro factors (e.g. wind, temperature, demand profiles) do not materially shift flows over a single 10-minute interval, meaning short-interval variability is dominated by stochastic volatility rather than structural change.

Under this interpretation, the distribution in Figure 9 represents the likelihood that actual metered flow will depart from the unbiased forecast by a particular amount due to this inherent volatility.


Although the precise forecast value at any moment is unobserved, a statistically implied safety margin, described in the [Reduced Requirement for Constraint Actions](#) section, can be estimated using the following assumptions:

- That the Control Room takes flow-reducing measures when forecasts indicate that flows will be close to, or exceed, the boundary limit.
- That these measures will seek to reduce the boundary flow to a level where the risk of exceedance is reduced.
- The historical rate of exceedance at a boundary at times when constraint actions were taken gives an indication of willingness to accept an exceedance.

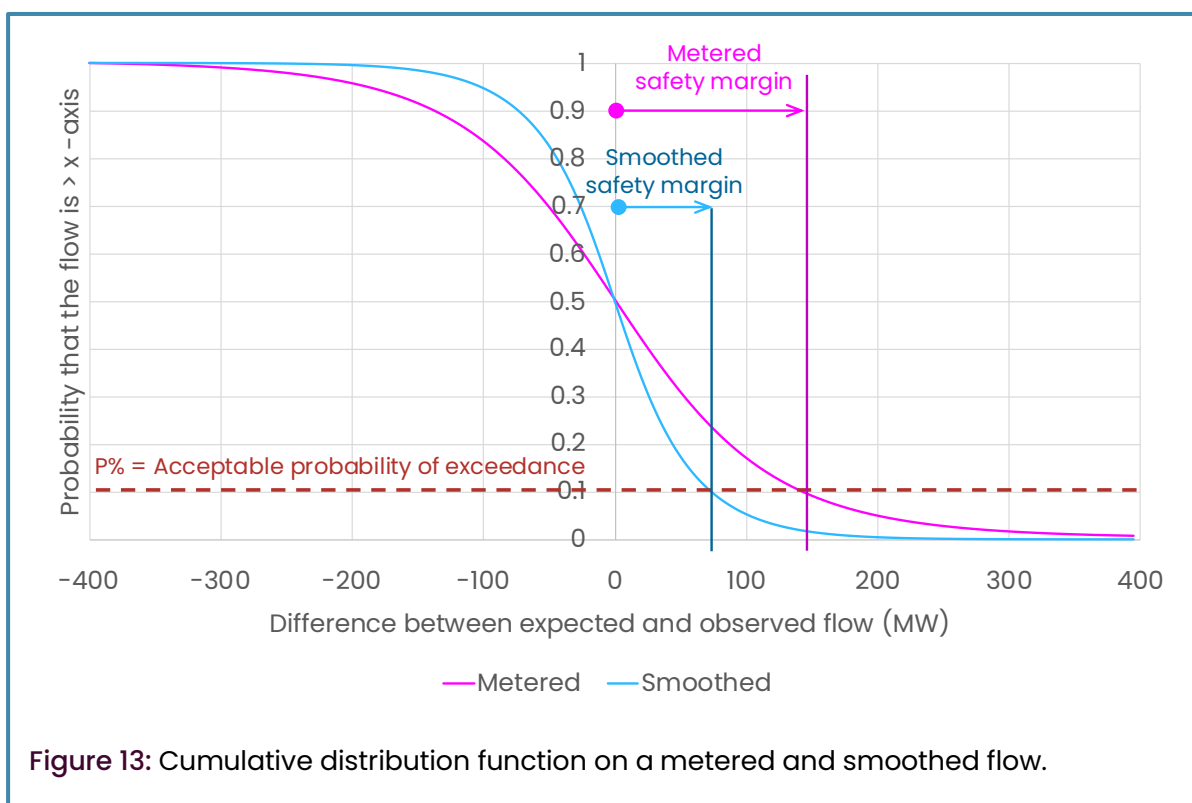
Implied Safety Margin Calculation

Given the basis of approach described above, the first difference analysis method is applied as follows:

1. Evaluate the distribution of first differences according to the method described in the [10-Minute Averages – Evolution of the Mean](#) section for both the smoothed and metered power flow data. The analysis includes only the highest 30% of flow periods and only the data where both periods were unconstrained.
2. Use the metered data to determine the percentage of the time, P , that the constraint limit is exceeded. This quantity gives an informal estimate of how 'acceptable' exceedance is.
3. Assume that BOA decisions are taken with the intention of remaining below the limit with a probability P of exceeding, then the P th centile of the first difference distribution is the size of the implied safety margin. This is illustrated in Figure 11.
4. The difference between the P th centile of the metered flow and the smoothed flow gives the reduction in implied safety margin that is consistent with maintaining historical precedent. This is additional flow that could pass through the boundary



without increasing the risk of exceedance. Figure 13 shows the difference in implied safety margin between a metered and smoothed flow. This difference is the amount of BOAs that can be unwound.



Although smoothing could be undertaken on any boundary at any time and would realise benefits due to lower peak flow levels, the potential to benefit from increases in flow due to a smaller implied safety margin can only be exploited when there are bids to the north and no constraints between the unit that was bid off and SSHARN3. This is tackled in the Nested Boundaries Analysis.


Nested Boundaries Analysis (3)

Inputs: Historical metered flows for all boundaries; Smaller implied safety margin; BM constraint-tagged BOA data.

Outputs: A revised set of flows with bids removed in line with lower implied safety margin calculations.

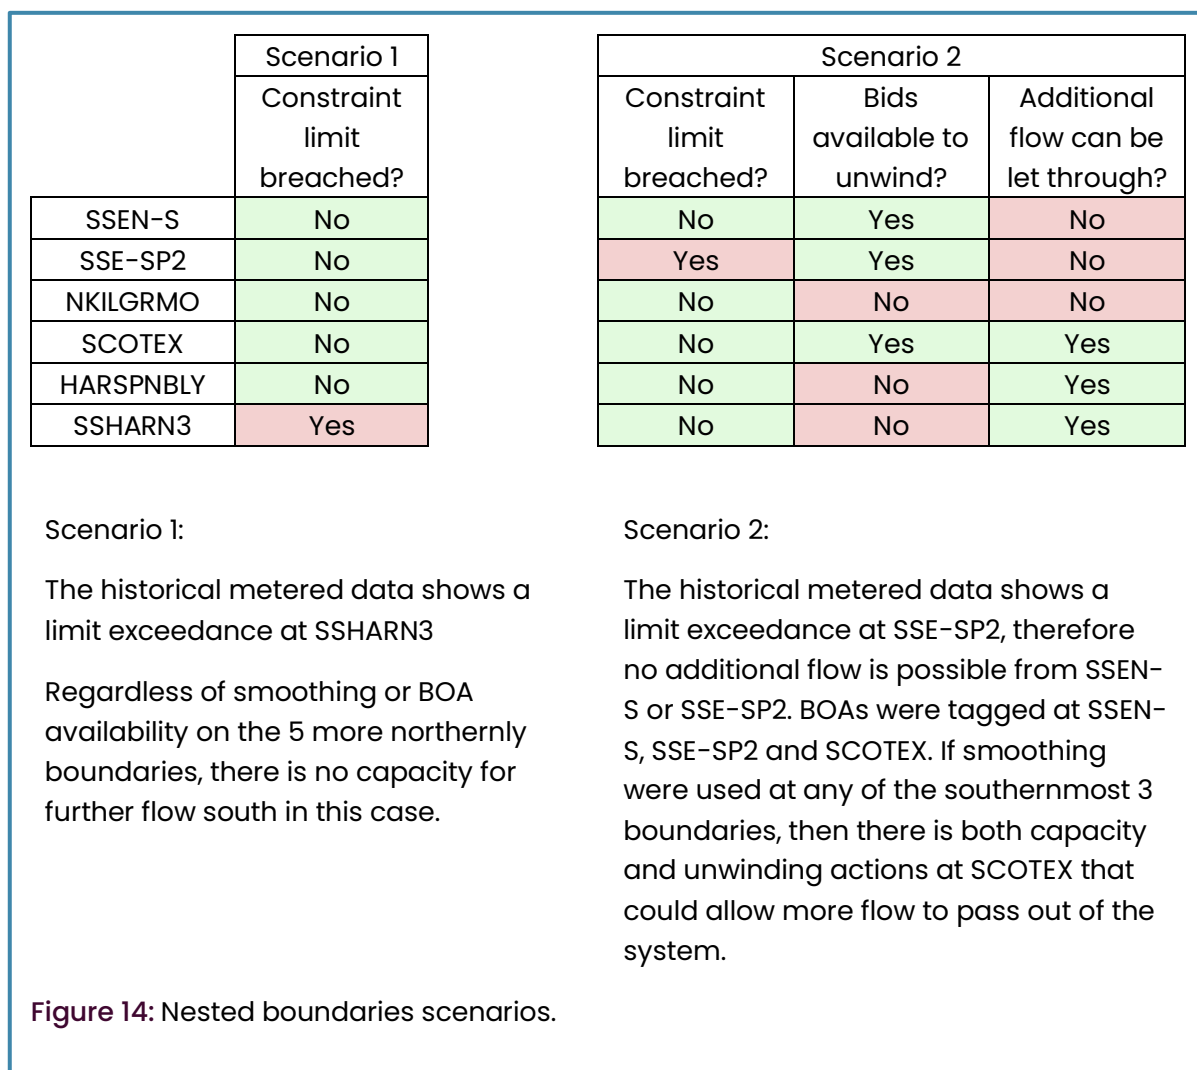
Nested boundaries introduce a significant complication to the process of quantifying the benefit of allowing more flow across boundaries. Some key features of the problem are:

1. Additional flow is only useful if it can leave the nested system, flowing southwards out of SSHARN3.
2. Additional flow needs to be available and accessible – there needs to be a BOA available to unwind to the north of the boundary where smoothing is done and, for



that energy to exit the system, there must be no constraints between the BOA and SSHARN3.

This is illustrated in two schematic examples in Figure 14.




From Scenario 1:

- The limit exceedance at SSHARN3 precludes any additional flow from flowing out of the region. Regardless of what boundary capacities exist to the north, smoothing cannot be used to justify a reduction in implied safety margin and increased flow, since any unwound BOAs would simply be reinstated in additional constraint actions required to manage flow at SSHARN3.

For Scenario 2:

- There is a limit exceedance at SSE-SP2. For the reasons given for Scenario 1, this means that the BOA actions tagged to SSEN-S and SSE-SP2 cannot be unwound and there is no capacity for additional energy to flow south at SSE-SP2.
- There are no BOAs available for unwinding at NKILGRMO, so the spare capacity there cannot be used.



- At SCOTEX, however, there is both spare capacity and available BOA actions that could be unwound. Smoothing the flow at any of SCOTEX, HARSPNBLY or SSHARN3 could potentially allow greater flow through a reduced safety margin. The additional flow is limited by whichever is the smallest of: the available implied safety margin, the available BOAs and the additional capacity that smoothing can realise.

The result is that the decisions affecting flows, smoothing and constraint management need to be considered across the system as a whole, yet this becomes computationally restrictive. To simplify this, the following approach is adopted.

- The system of metered flows is analysed every 10–minutes to assess where, if anywhere, within the nested set the following conditions for allowing additional flow have been met:

- Is the flow close to but not above the limit?


There is no procedural definition of 'close' that can be applied, therefore the analysis is repeated for three different definitions.

Where flows are further below the limit, no adjustment is made. If BOAs were used, it is assumed that there were reasons other than constraint limits for taking this action. Evidently, results will depend on where this threshold is applied, and sensitivity to this assumption is tested. No additional flow is allowed where the flow is already exceeding the flow limit.

- Is there capacity to the south?
- Are there bids and corresponding offers that can be unwound to provide additional energy flows?

Eligible bids are selected for removal based first on proximity to the boundary at which the implied safety margin is being reduced.

- If the three conditions are met, then the metered flows will be revised across the nested set. The size of the change will be the smallest of the spare flow capacity, the BOA volumes and the size of reduction in safety margin.
- The outputs of the nested boundary modelling are:
 - A revised set of self-consistent boundary flows based on metered flow data.
 - Changes to the historical BOA energy volumes at effected boundaries.
 - The net gain in flow out of the system i.e. the additional flow through SSHARN3.



BOA Cost-Saving Analysis (4)

Inputs: A revised set of flows for all boundaries; List of bid volumes unwound for every 10-minutes for each boundary; BM data on BOA volumes and costs; Carbon intensities by generator-type.

Outputs: Reductions in costs of constraint actions; Reduction in carbon emissions from constraint actions; Energy and power volumes needed to provide smoothing.

The final steps of the analysis calculate the monetary savings from reduced constraint actions, and the 'cost' incurred to provide smoothing, this being the energy and power flows required by the smoothing asset.


The nested flows analysis provides a list of the volumes of tagged bids at each boundary in each 10-minute period that could be unwound as a result of reducing the implied safety margin. These are cross referenced with the BM BOA data to estimate a cost saving. Where multiple bids are available at a boundary, options are unwound in descending cost order. All offers south of SSHARN3 are pooled and matched to unwound bids, again in order of descending cost.

BM bid data is available as energy reductions over each half-hour settlement period. To derive the additional power available over 10-minutes it is assumed that the units in question were bid-off for the whole settlement period; hence the power available in each 10-minute in MW is the energy (MWh) multiplied by 2.

Carbon savings may be estimated from the BM BOA data, as generator type is listed along with BOA costs and volumes. Assumptions on the carbon intensity of each generator type are listed in Table 2.

Table 2: Carbon intensity values used for BM generator types

Generator types (from BM data)	CO2 emissions factors (kg CO2e/MWh)	Source
Aggregated	0	Assume 0 as can be aggregated low carbon technologies or demand assets
Battery	0	Assume 0 as uses zero carbon energy
Biomass	120	NESO [1]
Combined Cycle Gas Turbine (CCGT)	394	NESO [1]
Combined Heat and Power (CHP)	610	BEIS [2]
Coal	937	NESO [1]



DIESEL	935	NESO - used OIL [1]
GAS	394	Assumed CCGT as no NESO definition in [1]
INTERCONNECTOR	337	Mean value of reported carbon intensity [3]
Non-Pump Storage Hydropower (NPSHYD)	0	NESO [1]
Open Cycle Gas Turbine (OCGT)	651	NESO [1]
PUMP STORAGE	0	NESO [1]
WIND	0	NESO [1]

The power and energy flows required to deliver smoothing that would give rise to the above savings were found by smoothing the revised boundary flows from the **Nested Boundaries Analysis (3)** section and outputting the power flows from the smoothing asset. This mimics a process in which:

- Forecast data and volatility concerns are combined to give a new set of BOAs.
- These result in a revised set of boundary flows.
- These revised boundary flows drive the actions of the smoothing algorithm.

The smoothing algorithm is applied to operate during the defined constrained periods.


Cost Benefit Analysis

The costs of providing a smoothing service include the following upfront elements:

- The cost associated with installing a suitable low-latency feed from where the boundary flow is monitored to the asset(s) that provides smoothing.
- The cost of developing a robust algorithm that is optimised for the type of asset, characteristics of the deployment location, and designed not to interfere with grid-stability related interventions (e.g. frequency response services).
- The cost of developing, consulting on and establishing a new market.
- The cost of implementing the service in control room systems to enable service deployment and monitoring.

The operational costs, incurred on an on-going basis include:

- Administration of the market.
- Paying service providers, with the basis of remuneration being unknown at this stage, but likely to be some combination of availability (paid per hour or settlement



period); availability incorporating a power consideration (paid per MW per hour); energy throughput (paid per MWh).

- Maintaining the service deployment, communication and monitoring infrastructure.

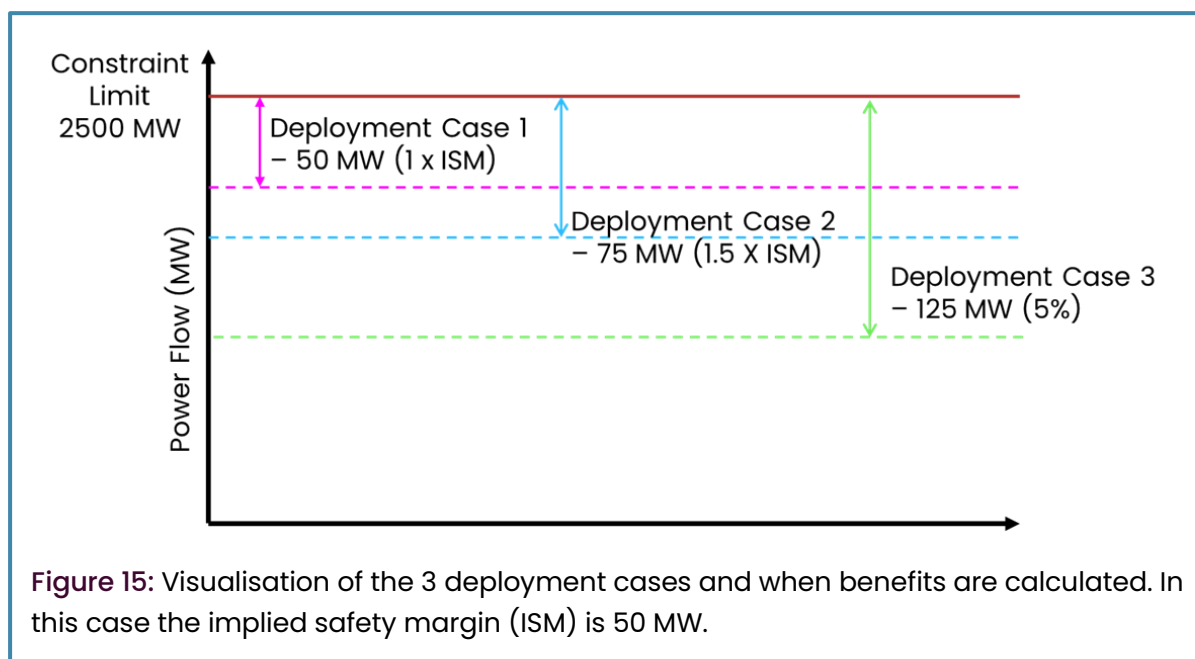
For boundary flow smoothing to make economic sense, the savings through reductions in constraint costs must be significantly greater than the costs listed above, over the lifetime of the intervention. There is also the need to consider the opportunity costs of the service providers: What could they earn instead, if they were not providing smoothing?


At this stage, there is no information on any of the upfront costs, nor on the likely costs of providing administration for the market. Therefore, the CBA does not calculate the total cost of providing a smoothing service. Instead, it takes the energy and power required to provide smoothing and uses this to calculate a ceiling price per MWh and per MW for the service. These values will be compared with the earnings available to possible providers of smoothing to understand whether this would be attractive to them.

Three operational margin deployment cases are considered for when smoothing would be deployed to enable a potential reduction in constraint actions:

1. Deployment Case 1: Additional flow is let through when the flow is within the implied safety margin of the flow limit.
2. Deployment Case 2: Additional flow is let through when the flow is within 1.5 x the metered implied safety limit.
3. Deployment Case 3: Additional flow is let through when the flow is within 5% of the constraint limit.

The three deployment cases are visualised in Figure 15.






This section has described the method used to determine the cost-benefit and the different scenarios being investigated. The results of the quantified analysis are presented in the following section.

5. Cost Benefit Analysis

Results





This section presents the results of the analysis in 5 tables, along with points to note.

Table 3: Historical metered flow likelihood of exceedance given that constraint actions were taken⁴, energy transferred over the limit for unsmoothed flows and implied safety margin calculations for smoothed and unsmoothed flows. Outputs based on metered data collected between 01/01/2024 and 30/03/2025 (15 months). Smoothing was carried out using the EWMA algorithm.

		SSEN_S	SSE_SP2	NKIL'MO	SCOTEX	HARS'BLY	SSHARN3
Historical exceedance probability, given constraint actions taken	%	7.9	20.7	0.8	16.0	24.7	6.9
Historical energy transferred over the limit	MWh	11,247	78,413	1,491	412,279	308,446	16,117
Implied safety margin for unsmoothed, metered flows	MW	92.5	54.6	236.9	95.3	66.0	173.6
Implied safety margin for smoothed, metered flows	MW	72.8	44.9	159.8	83.3	61.7	152.7
Reduction in safety margin due to smoothing	MW	19.7	9.7	77.1	12.0	4.2	20.9

Points to note from Table 3:

- The implied safety margin is inversely proportional to historical exceedance probability.
- The reduction in safety margin introduced by smoothing is also inversely related to the historical exceedance probability.

⁴ Constraint status determined as being when the adjusted flows exceed the constraint limit.




Table 4: Results of the nested boundaries analysis calculating the overall benefits in increased flows through SSHARN3 and reduced BOA actions from deploying the three different operational margin deployment cases for when smoothing benefits can be deployed. Outputs are based on metered data collected between 01/01/2024 and 30/03/2025 (15 months). Smoothing was carried out using the EWMA algorithm. Results show cost savings but do not include the costs of establishing and running a market for smoothing services.

		Deployment Case 1	Deployment Case 2	Deployment Case 3
Additional boundary flow	MWh	6,247	31,739	34,536
Hours of increased flow	hrs	361	1,176	2,366
Average increased flow	MW	17.3	27.0	14.6
BM Cost savings over 15 months	GBP	888,000	4,374,000	5,051,000
BM Cost savings (annual estimated)	GBP/yr	710,400	3,499,200	4,040,800
CO2 saved over period	tonnes	1,520	7,600	7,530
CO2 saved (annual)	tonnes/yr	1,220	6,080	6,030

Points to note in Table 4:

- Results show high sensitivity to the deployment assumptions.
- Monetary savings are of the order of single millions of pounds, before accounting for the costs of establishing and delivering the service.





Table 5: Results of analysis of algorithm flows at each boundary, including estimated power and energy storage requirements for the assets providing the smoothing. The results apply to each of the different sets of assessment assumptions, as algorithm flows are almost independent of the decisions on when to unwind BOAs. Outputs are based on metered data collected between 01/01/2024 and 30/03/2025 (15 months). Smoothing was carried out using the EWMA algorithm.

		SSEN-S	SSE-SP2	NKILGRMO	SCOTEX	HARSPNBLY	SSHARN3	TOTAL
Hours smoothing assets in use	Hours	4,600	6,050	4,360	3,480	3,570	3,200	25,260
Total flow to smoothing assets	MWh	89,600	132,800	94,800	126,800	134,000	132,600	710,600
Average power flow to/from assets	MW	19.5	21.9	21.7	36.5	37.5	41.4	-
Max 10-minute average power flow to/from assets	MW	379	460	393	601	579	622	-
Peak power flow to assets	MW	2,220	1,930	590	2,900	1,800	890	-
Total energy storage capacity	MWh	314	295	339	814	743	825	3,300
Exceedance probability of revised flows	%	6.0	19.1	0.6	15.0	23.6	5.8	-
Energy transferred over the limit for revised flows	MWh	7,500	67,274	1,107	409,205	303,263	12,791	801,140
	% historical ⁵	66.7	85.8	74.3	99.3	98.3	79.4	96.8

⁵ Historical values may be found in Table 3.



Points to note in Table 5:

- The condition that smoothing should operate in half-hour periods (aligned with BM settlement periods) when flows have been constrained, plus a half-hour before and a half-hour after, leads to high utilisation of smoothing.
- Peak power flows to the smoothing assets may be very large.
- Both exceedance probability and estimated energy transferred over the limit are close to the values for the historical metered flow. This is a desirable outcome as it indicates that most of the smoothing benefit has been captured by a decrease in safety margin.
- To achieve smoothing, the algorithm requires very significant volumes of energy available and high power requirements.






Table 6: Results from the nested boundaries analysis for the three sets of operational margin deployment cases. [Additional flow] is the additional flow that was able to pass through each boundary following application of smoothing to all boundaries. [Tagged bids unwound] is the volume of unwound bids at each boundary, and their value [Value of Bids unwound]. [Increase in energy transferred following safety margin change] shows the additional energy flow that resulted from the adjustment in safety margin at each boundary, and the value that this realised [Value of safety margin change]. Outputs are based on metered data collected between 01/01/2024 and 30/03/2025, with annualised values also presented. Smoothing was carried out using the EWMA algorithm. Results show cost savings but do not include the costs of establishing and running a market for smoothing services.

DEPLOYMENT CASE 1								
		SSEN-S	SSE-SP2	NKILGRMO	SCOTEX	HARSPNBLY	SSHARN3	TOTAL
Additional flow across boundary	MWh	820	3,300	5,930	6,010	6,060	6,250	-
Total energy flow to smoothing asset	GWh	89.6	132.8	94.8	126.8	143.0	132.5	710.6
Tagged bids unwound	MWh	820	2,480	2,630	80	50	190	6,250
Value of Bids unwound @ boundary	£	136,000	357,000	338,000	15,000	9,000	33,000	888,000
	£/yr	112,400	295,200	279,500	12,400	7,400	27,300	734,200
Increase in energy transferred following safety margin change ⁶	MWh	700	360	4,810	60	20	300	6,250
Value attributable to change in safety margin @ boundary	£	117,000	63,000	649,000	10,000	3,000	46,000	888,000
	£/yr	96,700	52,100	536,600	8,300	2,500	38,000	734,200

⁶ A reduction in the implied safety margin at a boundary allows greater flow of energy. This energy does not necessarily come from the unwinding of a bid at that boundary – it could come from the unwinding of a bid further north. Hence, for example, at SSE-SP2, 2,480 MWh of bids were unwound, but this principally is a result of safety margin reductions further south: only 360 MWh of additional flow arose as a result of smoothing and safety margin reductions at SSE-SP2 itself.




DEPLOYMENT CASE 2								
		SSEN-S	SSE-SP2	NKILGRMO	SCOTEX	HARSPNBLY	SSHARN3	TOTAL
Additional flow across boundary	MWh	4,940	17,650	28,610	29,410	30,000	31,740	-
Total energy flow to smoothing asset	GWh	89.6	132.6	95.4	126.7	133.8	132.3	710.6
Tagged bids unwound	MWh	4,940	12,710	10,960	800	590	1,740	31,740
Value of Bids unwound @ boundary	£	759,000	1,714,000	1,372,000	136,000	108,000	285,000	4,374,000
	£/yr	627,500	1,417,100	1,134,400	112,400	89,300	235,600	3,616,400
Increase in energy transferred following safety margin change	MWh	3,890	2,540	21,640	500	340	2,830	31,740
Value attributable to change in safety margin @ boundary	£	632,000	467,000	2,635,000	102,000	77,000	461,000	4,374,000
	£/yr	522,500	386,600	2,178,700	84,300	63,700	381,400	3,617,200



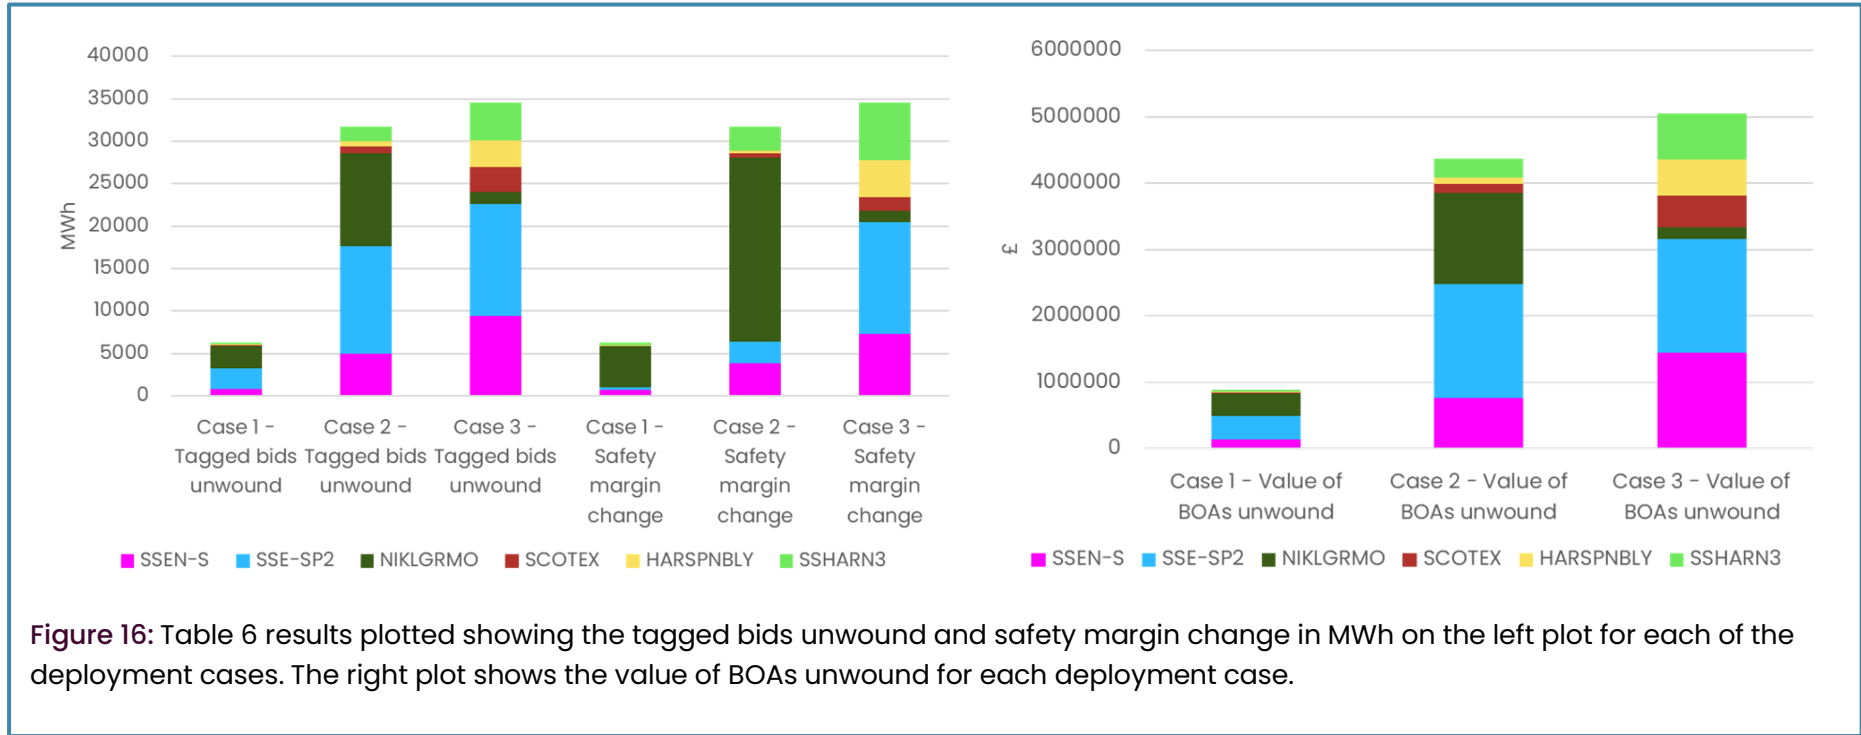
DEPLOYMENT CASE 3								
		SSEN-S	SSE-SP2	NKILGRMO	SCOTEX	HARSPNBLY	SSHARN3	TOTAL
Additional flow across boundary	MWh	9,500	22,600	24,100	26,900	30,100	34,500	-
Total energy flow to smoothing asset	GWh	89.1	131.5	94.1	126.2	133.5	132.0	706.3
Tagged bids unwound	MWh	9,500	13,100	1,500	2,900	3,200	4,400	34,600
Value of Bids unwound @ boundary	£	1,444,000	1,716,000	173,000	475,000	548,000	695,000	5,051,000
	£/yr	1,193,900	1,418,800	143,000	392,700	453,100	574,600	4,176,200
Increase in energy transferred following safety margin change	MWh	7,330	13,160	1,310	1,610	4,430	6,760	34,600
Value attributable to change in safety margin @ boundary	£	1,117,600	1,798,600	201,000	261,300	680,900	991,600	5,051,000
	£/yr	924,000	1,487,100	166,200	216,100	563,000	819,800	4,176,200

Points to note in Table 6:

- The boundary at which the additional capacity was created and the boundary at which BOAs were unwound to create the extra flow are often different. This is why a difference in value is seen between [*Value of BOAs unwound*] and [*Value of safety margin change*].
- The influence of smoothing at a boundary, in terms of realising additional flow capacity, is dependent on the criteria for choosing when a reduction in safety margin may be applied. This is exemplified at NKILGRMO: Under Deployment Cases 1 and 2 a reduction in safety margin is implemented when the flow is within a distance of the flow limit that is determined by the size of the safety margin. The implied safety margin at NKILGRMO is large (due to a low probability of exceedance) so the condition for reducing the safety margin is often met and NKILGRMO appears to be the best boundary for smoothing. However, Deployment Case 3 allows safety margin reduction within a given % of the flow limit. This dramatically reduces the occasions on which additional flow is permitted at NKILGRMO, and it becomes the least favourable boundary for smoothing.



- The small discrepancy between the total values calculated according to where BOAs were unwound and according to where the original flow increase was created is due to the assumption that BOA power flows are constant across a settlement period.








Table 7: The value arising from changing the safety margin at each boundary (Table 6 row 6) are divided by the time, power and energy used to deliver the service for each of the deployment assumptions to give per unit savings. Results show savings but do not include the costs of establishing and running a market for smoothing services. They represent an absolute upper bound on the price that NESO should be willing to pay for smoothing.

DEPLOYMENT CASE 1								
Average saving in cost of BOAs per ...	Unit	SSEN-S	SSE-SP2	NKILGRMO	SCOTEX	HARSPNBLY	SSHARN3	WEIGHTED AVERAGE
... hour of smoothing	£/hr	25.35	10.48	148.53	2.84	0.81	14.47	35.10
... MWh of energy transferred to smoothing asset	£/MWh	1.30	0.48	6.83	0.08	0.02	0.35	1.25
... MW of smoothing asset capacity required per year (overall peak)	£/MW/yr	44.00	27.00	909.00	3.00	1.00	43.00	71.00
... MW of smoothing asset capacity required per year (peak 10-min mean)	£/MW/yr	255.00	113.00	1,367.00	14.00	4.00	61.00	242.00
... MWh of asset storage per year	£/MWh/yr	307.00	178.00	1,579.00	10.00	3.00	46.00	220.00



DEPLOYMENT CASE 2								
Average saving in cost of BOAs per ...	Unit	SSEN-S	SSE-SP2	NKILGRMO	SCOTEX	HARSPNBLY	SSHARN3	WEIGHTED AVERAGE
... hour of smoothing	£/hr	137.40	77.30	604.40	29.30	21.60	144.10	173.20
... MWh of energy transferred to smoothing asset	£/MWh	7.05	3.52	27.80	0.80	0.57	3.48	6.16
... MW of smoothing asset capacity required per year (overall peak)	£/MW/yr	235.00	200.00	3,693.00	29.00	35.00	429.00	350.00
... MW of smoothing asset capacity required per year (peak 10-min mean)	£/MW/yr	1,378.00	841.00	5,550.00	140.00	110.00	613.00	1,192.00
... MWh of asset storage per year	£/MWh/yr	1,664.00	1,311.00	6,427.00	104.00	86.00	462.00	1,086.00



DEPLOYMENT CASE 3								
Average saving in cost of BOAs per ...	Units	SSEN-S	SSE-SP2	NKILGRMO	SCOTEX	HARSPNBLY	SSHARN3	WEIGHTED AVERAGE
... hour of smoothing	£/hr	243.00	297.00	46.00	75.00	191.00	310.00	200.00
... MWh of energy transferred to smoothing asset	£/MWh	12.50	13.50	2.10	2.10	5.10	7.50	7.10
... MW of smoothing asset capacity required per year (overall peak)	£/MW/yr	416.00	771.00	282.00	75.00	313.00	921.00	404.00
... MW of smoothing asset capacity required per year (peak 10-min mean)	£/MW/yr	2,437.00	3,234.00	423.00	359.00	972.00	1,317.00	1,377.00
... MWh of asset storage per year	£/MWh/yr	2,943.00	5,041.00	490.00	265.00	758.00	994.00	1,254.00

6. Discussion and Conclusions

Overview of Findings and Sensitivity Assessment*

Potential for Constraint Cost Reduction


Comparison of Costs and Potential Benefits

Technology Suitability and Opportunity Cost

Summary Conclusions

Suggested Further Work





Overview of Findings and Sensitivity Assessment

This section discusses the findings of the quantitative analysis of the constraint cost reductions and examines the extent to which they depend on the underlying assumptions. It evaluates the classes of assets capable of delivering the proposed service and considers the opportunity costs faced by providers when offering boundary flow smoothing. The section also identifies areas where additional investigation would materially improve confidence in the assessment. Finally, it provides conclusions regarding the feasibility of implementing boundary flow smoothing.

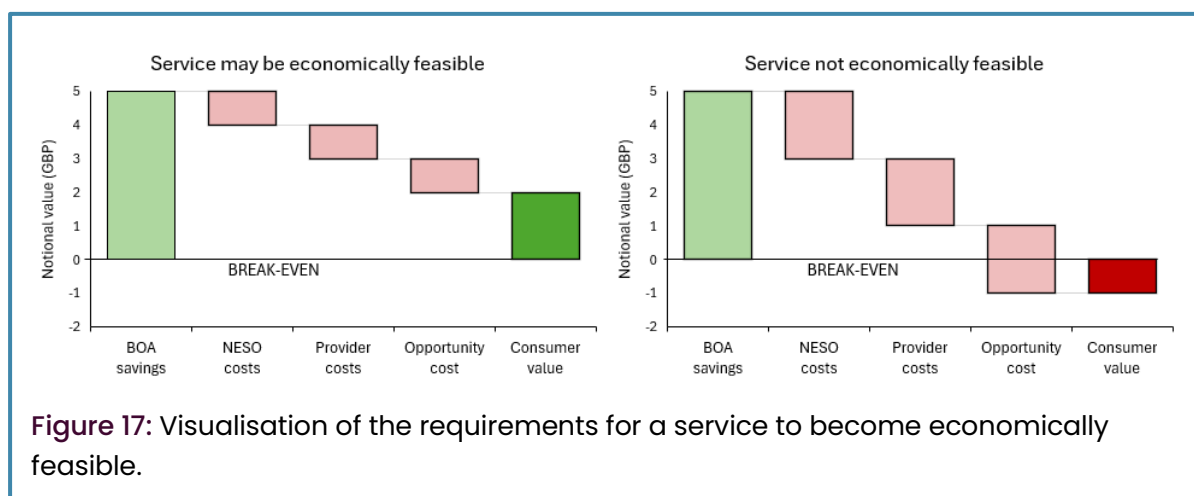
A full list of project assumptions may be found in Appendix M – Project Assumptions.

Conditions for Economic Feasibility


For boundary flow smoothing to be economically worthwhile, the service must create sufficient value to cover the costs faced by all parties. In practical terms, the reduction in constraint costs needs to be large enough to:

- Meet NESO's costs of developing, operating, and administering the service.
- Offer service providers a return that is more attractive than their next-best alternative use of the asset, and cover their participation costs.
- Still leave a meaningful benefit for consumers.

This is illustrated in Figure 17. A further condition is that, even if it yields a benefit for consumers that is greater than the costs of implementation, there is still the question of whether NESO's implementation costs could provide greater benefit if invested elsewhere.



These requirements translate into three conditions necessary for boundary flow smoothing to fulfil in order to achieve economic feasibility:



- It must reduce the costs of constraint-related BM actions (BOAs).
- Any savings must exceed the full cost of delivering the service, including low-latency communications, market establishment, ongoing operational costs and paying service providers.
- The ceiling price implied by the net savings must be sufficient to attract market participants, and cover their participation costs.

The extent to which these requirements have been met will be addressed in the following sections.

A caveat


The validity of all the findings is predicated on the assumption that the Control Room is factoring volatility into its decisions on flow management.

Unlike peak shaving or demand response, smoothing the flow impacts only the magnitude of fluctuations and does not change the mean of the flow over settlement period timescales. This service can therefore only lead to additional boundary flow if volatility reduction prompts changes to Control Room decisions on BOA actions. To assess the impact of smoothing therefore requires an understanding of how, if at all, volatility is currently accommodated in decisions (explicitly or implicitly).

In the past, a thermal safety margin was deployed in some instances to reduce the probability that short-term volatility would cause exceedances. This safety margin represented an operational gap between the boundary constraint limit and the operational flow level to be targeted by the Control Room. At present however, there is no formal procedure that defines a safety margin in the context of thermal boundary flow constraints.

To progress to an evaluation of boundary flow smoothing in the absence of a formal Control Room procedure, this work has developed an approach that derives an 'implied' safety margin. This uses the volatility behaviour of the observed boundary flows to infer a flow level to be targeted for each boundary that is statistically consistent with observed exceedance probabilities across a large sample of 10-minute intervals. Repeating this process using smoothed historical flows delivers a second, smaller implied safety margin, with the difference between the two indicating a potential for increased flow in the smoothed case.

However, if the Control Room does not incorporate volatility into operational decisions and routinely targets the thermal flow limit irrespective of the likelihood of exceedances, then smoothing will not reduce BOA volumes. In this case, any benefits of smoothing would be limited to reduced mechanical or thermal stress on transmission assets, rather than reduced constraint payments.



Potential for Constraint Cost Reduction

The results indicate that reduction in the use of BOAs would lower constraint spending by between £0.7 million and £4 million pounds per year (Table 6). The range is wide because cost savings are highly sensitive to the assumptions used to infer when an implied safety margin might have been applied and when smoothing would plausibly have allowed that margin to be reduced.

Interpretation of Estimated BOA Reductions

Influence of Boundary Operating Practices

In Table 3, the estimated benefits at each boundary are seen to correlate with the observed exceedance probability. A high exceedance frequency will generate a small implicit safety margin, consistent with the Control Room applying fewer BOAs at that boundary and targeting higher flows. A smaller implicit safety margin reduces the scope for smoothing to change operational decisions and therefore limits the achievable benefits.

Sensitivity to Deployment Criteria

This analysis contains no assumptions on operational deployment strategies for a flow smoothing service. Table 4 demonstrates that the estimated benefits are highly sensitive to the deployment criteria used to determine when a safety margin reduction could be applied i.e. how close to the constraint limit the flow is. This sensitivity is to be expected, since benefits scale approximately with the number of settlement intervals over which savings are assumed to occur. The three deployment cases produce benefits of a similar order of magnitude, and the analysis highlights the need to specify operational decision rules explicitly when developing future deployment models.


Effects of Incomplete Constraint Data

The absence of a complete and accurate record of constraint-tagged BOAs introduces uncertainty into the estimated potential benefits of smoothing. The analysis relies on the set of actions currently tagged as constraint-driven, but retrospective tagging is known to omit some actions and misclassify others. As a result, the true scale of constraint-related activity is uncertain, and the available dataset may not fully represent the actions that smoothing could help address.

Assumptions Governing Safety Margin Calculation

Several assumptions used in calculating the implicit safety margins are likely to lead to over-estimation of the potential benefits of smoothing.

First, the analysis attributes all variation between successive 10-minute flow values to short-term noise. In practice, part of this variation reflects underlying macro changes, such as time-of-day and broader weather conditions, rather than true volatility. This



inflates the inferred spread of first difference flows and consequently exaggerates the amount of volatility that smoothing could remove. Access to forecast flows or methods that filter out lower-frequency trends would allow a more accurate separation of noise from underlying changes.

Second, the modelling assumes an instantaneous response from the smoothing asset. Real-world implementation will involve communication delays and finite ramp rates. These factors mean assets would respond to slightly outdated flow information and would not fully cancel short-term fluctuations. The resulting residual volatility would reduce the feasible reduction in implicit safety margins.

Third, it is also assumed that the asset that is providing the smoothing will operate at 100% efficiency. Losses due to round-trip inefficiencies of assets will have an impact on the algorithm behaviour and introduce an energy penalty to the smoothing process.

Fourth, it is assumed that all the power from the asset reaches the boundary. In practice, network losses mean that the flows to or from the asset would be less than expected, reducing the impact of the smoothing.

Together, these assumptions lead to an optimistic estimate of both the volatility available to be mitigated and the effectiveness of smoothing in practice.

Implications for Overall Benefit Magnitude

On balance, the benefits of smoothing are likely to be over-estimated in this analysis. This primarily reflects assumptions that overstate both the volatility that smoothing can remove and the effectiveness with which a real asset could respond. The first difference approach likely attributes too much variation to short-term noise, and the simulations assume instantaneous response, which is not operationally realistic. In addition, the benefits are highly sensitive to assumptions about when a reduced implied safety margin would be applied, and there is no established operational practice to anchor those assumptions. Collectively, these factors increase the likelihood that the analysis overstates the reduction in constraint actions achievable through smoothing.


Conclusion 1:

Indicative reductions in BOA constraint payments totalling up to the-order-of single millions of pounds per year are possible, though it is likely that this analysis over-estimates the savings.

Energy and Power Requirements for Smoothing Operation

The energy and power transfers required to deliver smoothing, the implied storage capacities, and the operational durations are large (Table 5). Normalising the BOA reduction savings to account for asset time, energy and power gives the following, in which the initial figure is the weighted mean and the bracketed figures give the P10-P90 range:

- £136 per hour of smoothing activity (£10 - £300);



- £4.80 per MWh of energy moved to provide a smoothed flow (£0.30 – £12.50);
- £275 per MW of power capacity (peak) for the smoothing asset (£27 – £910);
- £937 per MW of power capacity (10-min mean) for the smoothing asset (£61 – £2440);
- £220 per MWh of energy storage capacity for the smoothing asset (£46 – £2940).

This section discusses the operational requirements of boundary flow smoothing, focusing on the energy throughput, peak power, and implied storage needed. It also explains why these requirements are comparatively insensitive to some of the benefit-related assumptions and identifies how the deployment logic assumed here may inflate operational volumes.

Summary of Estimated Operational Requirements

Across the analysed period, smoothing requires substantial energy throughput and high peak power. This arises from the objective of counteracting short-term boundary flow volatility, which can require frequent changes in import and export power. As a result, the service is operationally intensive even when the associated monetised benefits are modest.

The required power and energy capacities remain key determinants of technical feasibility and of the likely provider set, since they govern asset eligibility, cycling burden, and procurement volume.


Insensitivity to Smoothing Deployment Cases

Results in Table 6 show that the energy flows used in deploying smoothing are largely insensitive to the different operational margin deployment cases regarding when smoothing might enable a reduction in bids and an increase in flow. This reflects the following:

- That smoothing was active during constrained periods (defined in the [Historical Boundary Flow Behaviour](#) section), and for a half-hour before and a half-hour after, to accommodate the ramping of the algorithm. This is the same across all three cases.
- That the increases in flow that smoothing might allow do not materially change the work that the smoothing asset has to do. This is linked to the volatility of the flow that it receives, which is not significantly altered by the addition of an extra 10–20 MW of flow. Also noting that the added flows themselves have zero volatility in this idealised case.

Implications of Nested Constraints and Limited Actionable Intervals

From the perspective of reducing constraint-related BM actions, the fact that smoothing is active whenever there are constraints means that it inevitably occurs during intervals where it cannot yield a benefit in terms of BOA reduction. This is due to system bottlenecks (Figure 14), which mean that smoothing may be operationally active at a boundary even when restrictions elsewhere prevent any increase in transferable power or any meaningful change in BOA selection.



This highlights that smoothing does not alleviate constraints: It adds no physical capacity to the network and can only permit small increases in power flows as they approach the constraint limit. This is particularly acute within the nested boundaries considered here owing to the high likelihood of there being a constraint somewhere within the group.

Another way in which this analysis over-estimates the energy and power transfer required for smoothing is through operating smoothing on each boundary flow independently. This neglects any downstream pass-through of lower volatility flow, for example:

- Smoothing downstream would become less challenging.
- It might be possible to reduce the number of bid actions taken at a boundary on the basis of smoothing conducted upstream.

Refining when smoothing is deployed, so that it is not used when it cannot potentially lead to higher flow, would reduce smoothing run time and energy throughput. This would lower service delivery costs and increase estimates of the achievable provider revenue per unit of contracted capacity.

The interplay between smoothing actions taken across the nested set, and the impacts of this on the opportunity for BOA reduction is likely to be complex. Further investigation could be undertaken to ascertain whether much of the benefit calculated here could be achieved through focussing smoothing activity at particular boundaries.

However, such refinements would primarily reduce the cost of delivering the service rather than increase system savings. The upper bound on monetisable benefit remains constrained by the extent to which reduced volatility translates into different operational decisions and reduced BOA volumes.

Interpretation and Associated Conclusions

The findings indicate that operational requirements are dominated by the deployment logic and by the underlying volatility of boundary flows, rather than by assumptions about whether smoothing enables incremental increases in flow. This supports higher confidence in the power and energy capacity requirements than in the estimated financial savings, although both remain sensitive to how realistically asset dynamics and control delays are represented.


Conclusion 2:

Smoothing alone does not add capacity to the constrained system and, in a highly constrained nested system, it is likely that bottlenecks will restrict the ability of the Control Room to enhance flow through exploiting lower volatility.

Conclusion 3:

The volumes of energy moved and the time for which smoothing is operational are likely to be over-estimated in the current analysis.

Conclusion 4:



The power and energy capacities required of service providers are critical to service viability. These requirements are primarily determined by the smoothing algorithm rather than deployment decisions and are therefore assessed with higher confidence.

Comparison of Costs and Potential Benefits

As noted previously, smoothing is only worthwhile if its benefits exceed its costs. These costs include not only payments to service providers, but also the costs of establishing and operating the service, such as telemetry, low-latency communications, market systems, assurance, compliance, and operational processes.

This study is an early order-of-magnitude feasibility assessment, and these wider costs are not quantified. Nevertheless, it is likely that the combined costs of infrastructure, market establishment, and on-going operations could be in the order of single millions of pounds per year, which would be comparable to, and could potentially eliminate, the estimated savings.

Conclusion 5:

The costs of enabling infrastructure and of establishing and operating the service are likely to be similar in magnitude to the projected savings. It is therefore uncertain whether the benefits would outweigh the full costs.

Technology Suitability and Opportunity Cost

Technology Suitability

Technology archetypes were developed by first performing a suitability analysis across a broad set of flexibility technologies, then establishing typical asset capabilities relevant to smoothing through research and stakeholder engagement.

Wider feedback on whether a smoothing service would be attractive to asset owners and service requirements is presented in 'Appendix I – Stakeholder Engagement Feedback'. The technology options, suitability analysis, and archetypes are provided in 'Appendix J – Technology Options', 'Appendix K – Technology Options Suitability Analysis' and 'Appendix L – Technology Archetypes' respectively.

Table 8 shows the smoothing metrics for SCOTEX to assess asset requirements over the analysed 15-month period. SCOTEX was found to have the largest energy and power requirements.




Table 8: Smoothing asset requirements

Average Asset Import per Day	Average Asset Export per Day	Storage capacity required* (*assuming bidirectional asset)	Peak Import Power	Peak Export Power
420 MWh	394 MWh	814 MWh	2896 MW	777 MW

The capabilities important for smoothing are as follows:

- **Power rating** – it is assumed the service procurement will be made up from multiple assets so while assets with larger power ratings are beneficial, as it means fewer assets need to be contracted, any asset rating can theoretically provide smoothing. The localised nature of smoothing (assets need to be located north of a boundary) will mean that having large ratings on assets will be beneficial in order to meet service procurement requirements.
- **Energy capacity** – similarly to power rating, although large capacity is beneficial for procurement it does not mean that one technology should be prioritised over another.
- **Duration of sustained action** – constrained periods often last for long periods (several hours) therefore those assets that are able to sustain action for longer periods are more suitable assuming that the procurement periods align with longer constrained periods.
- **Power direction** – the service could procure export and import capacity separately so power direction does not have a large impact on technology preference and this means more types of technologies can participate if bidirectional behaviour is not required. However bidirectional behaviour is beneficial if the asset is energy limited as it means the asset can provide the service for longer. This is demonstrated by the storage capacity in Table 8 being lower than the average import and export values.
- **Response time** – the smoothing simulation has assumed that the algorithms affect the flows instantaneously. For real assets, this is unrealistic, as there will be some lag for the asset to receive the data feed and then respond to signals. Additionally, the smoothing must not interfere with frequency response which needs to be prioritised for system security. Therefore, the capability to react within seconds, slower than frequency response but much quicker than BM actions (so less are needed), is anticipated.
- **Ramp rate** – a quick ramp rate is beneficial for smoothing as it aims to reduce short term volatility. So if an asset takes several minutes to ramp up it won't be able to impact the short term volatility.

The capabilities shown in Table 9 are key to differentiate the applicable technologies: duration of sustained action, power direction, response time and ramp rate. These criteria have led to the following assessment of assets on their capabilities to provide smoothing.




Table 9: Key technology archetypes

Asset	Duration of sustained action (typical)	Power direction	Response time	Ramp rate
Battery Energy Storage Systems (BESS)	1 to 4 hrs (with trend of increasing duration)	Generation up/down Demand turn up/down	Sub-second (0.5 s baseline depends on communication and plant)	100% in <1 seconds
Long Duration Electricity Storage (LDES) BESS	8 to 10 hrs (>8 hrs end of life capacity requirement part of LDES cap and floor scheme)	Generation up/down Demand turn up/down	Sub-second (0.5 s baseline depends on communication and plant)	100% in <1 seconds
Flexible Data Centres	Unlimited	Demand turn up/down	<2 seconds	100% in <5 seconds
Small Hydro	Varies, subject to reservoir capacity	Generation up/down	Seconds-Minutes	100% in 2 minutes
Large Hydro	Varies, subject to reservoir capacity	Generation up/down Demand turn up/down (if pumped)	Minutes	100% in 3-5 minutes
Industrial Demand	Minutes-Hours	Demand turn up/down	Seconds-Minutes	100% / 3-8 minutes
Aggregated EV	Minutes-Hours	Demand turn up/down	Seconds-Minutes	100% / 3-8 minutes
District Heat	Minutes-Hours	Demand turn up/down	Minutes	100% in 10 minutes

The two most important capabilities are response time and ramp rate due to the requirement of boundary flow smoothing to manage short term volatility. The direction of power flows and time spent in each direction has been investigated using the available 15 months of data, as shown in Table 10.




Table 10: Power flow times and directions over 15 months.

Boundary	Average Import Time (s)	Average Export Time (s)	Export/Import Direction Changes
All	263	269	443,286
SSEN-S	438	453	44,058
SSE-SP2	433	453	44,287
NKILGRMO	228	237	84,414
SCOTEX	246	251	79,137
HARSPNBLY	398	404	48,974
SSHARN-3	138	137	142,416

On average power flows in one direction for around 4 minutes before changing directions. Therefore, assets that can respond quickly and ramp to 100% output quickly will be required, meaning that the following assets are most appropriate for smoothing:

1. BESS
2. LDES BESS
3. Flexible data centres

If the time between power flow changes increases in future then this may open up smoothing to the remaining assets, or if a slower acting smoothing capability was procured for longer periods of consistent behaviour (similar to frequency response or reserve services).

Conclusion 6:

BESS, LDES BESS and flexible data centres possess the technical characteristics that could make them suitable for delivering this service.

Opportunity Cost – BESS

Table 11 compares indicative BESS earnings (Modo Energy, [2]) with the per-unit maximum cost savings implied by the smoothing analysis (from Table 7). The comparison indicates whether a battery operator would plausibly forego existing revenue streams to supply smoothing.




Table 11: Opportunity cost comparison with historical BESS revenues


	Revenue per (peak) MW installed capacity / year	Revenue per (10-min mean) MW installed capacity / year	Revenue per MWh installed capacity / year
Historical data (01/01/24 - 31/03/25)	£57,026		£41,928
Deployment Case 1, no implementation or operations cost	£71 (£1 - £909)	£242 (£4 - £1367)	£220 (£3 - £1580)
Deployment Case 2, no implementation or operations cost	£350 (£29 - £3,693)	£1,192 (£110 - £5,550)	£1,086 (£85.7 - £6,430)
Deployment Case 3, no implementation or operations cost	£404 (£75 - £921)	£1377 (£359 - £6,427)	£1,254 (£265 - £5,041)

Once full delivery and participation costs are included, even the upper estimates are likely to reduce materially, to the order of hundreds of pounds per MW-year or MWh-year or less. At those levels, an asset owner would have limited incentive to provide smoothing unless alternative revenues were low or unavailable. Although BESS can stack revenue streams by simultaneously providing different services, smoothing is operationally intensive and may prohibit the ability to stack with other services. In addition, the high-frequency charge and discharge required for smoothing is likely to increase cycling and degradation, further reducing attractiveness unless explicitly compensated.

For flexible data centres, equivalent historical revenue benchmarks are not available. In addition, the algorithm may require modification because data centres often provide one-sided flexibility (for example, load reduction only). A more appropriate comparator is likely to be demand-side flexibility products. If compared on an energy-throughput basis, the implied smoothing value of approximately £1 to £7/MWh appears low relative to typical flexibility services, but this should be validated using a consistent product definition and relevant market benchmarks.

Conclusion 7:

Even with a likely over-estimate of the BOA savings and the exclusion of wider implementation and running costs, the revenues available from smoothing appear to be too low to be attractive for providers.



Summary Conclusions

▶ **Savings likely overstated:**

Smoothing could potentially enable more efficient use of existing transmission capacity by reducing the extent of volatility-related BOA actions, reducing annual BOA costs by up to the low single millions of pounds annually. These figures are assumption-sensitive and highly likely to be over-estimated.

▶ **Nested constraints limit value:**

Smoothing adds no physical capacity to the network. Where the system is already operating at capacity in places, such as in the nested set of boundaries considered here, the bottlenecks often prevent possible increased energy flows from unwound BOAs from propagating southwards.

▶ **Operational volumes over-estimated:**

Energy throughput and run time are likely inflated by the smoothing activation conditions and by modelling each boundary independently.

▶ **Capacity requirements are decisive:**

Required power and energy capacities drive feasibility and are more robustly estimated than monetary savings.

▶ **Net benefit uncertain:**

Wider infrastructure, market, and operational costs are considered likely to be comparable to or greater than the projected savings.

▶ **Candidate technologies:**

Fast response and ramping are essential; BESS, LDES BESS, and flexible data centres are the most suitable.

▶ **Low provider attractiveness:**


Implied revenues are low compared to opportunity costs and cycling burden, so provider participation is unlikely under the evaluated design.

▶ **Operational dependence:**

BOA reductions are only plausible if volatility is implicitly managed through safety margins or similar operational process, and if smoothing would change BOA selection. Otherwise, benefits are limited to reduced asset stress.

Overall, the feasibility case for boundary flow smoothing is weak against the three conditions for economic feasibility defined at the outset of this section and boundary flow smoothing does not present a robust business case for the nested boundary set assessed.

The analysis suggests potential BOA savings in the low single millions of pounds per year, but these are contingent on how the Control Room applies implied safety margins and



are likely overstated due to optimistic volatility and response assumptions. Once enabling and operational costs are included, net benefit is uncertain. The calculated ceiling price is low relative to provider opportunity costs and the operational intensity required, making participation unlikely. Taken together, and given that nested constraints frequently prevent smoothing from enabling additional transfer, smoothing is unlikely to be viable.

Suggested Further Work

The analysis undertaken in this study provides a comprehensive assessment of the technical feasibility, operational impact and economic case for boundary flow smoothing. The findings indicate that, while smoothing could theoretically reduce short-term volatility and marginally reduce constraint-related BOAs, the scale of achievable benefits is at best modest and highly sensitive to operational assumptions. When set against the energy and power requirements of the service, the operational intensity for providers, and the wider infrastructure and market costs needed to enable the service, smoothing does not present a compelling case for implementation across the boundaries examined.

However, if there is continued interest in exploring the concept further, either for specific boundaries, future system states, or targeted operational use, there are several areas where additional work would help refine understanding and confirm whether a niche application might still offer value.

Operational and Decision-Making Refinements

Although the current analysis reasonably represents operational behaviour, further work could deepen alignment with Control Room practice, particularly if smoothing were to be reconsidered in a scoped or experimental form.

Formal Operational Rules for Volatility Management


A more explicit model of how volatility, exceedance tolerance and BOA selection interact in real time would help clarify whether smoothing could ever materially influence operational decisions under evolving system conditions.

Integration of Forecast Data

Using actual short-term forecast or PNA predictive outputs would help separate forecast error from intrinsic volatility and indicate whether smoothing could help most during specific forecast-challenging periods (e.g. storm fronts, rapid wind ramps).

Targeted Network and Boundary Modelling

While smoothing appears unlikely to be beneficial when applied broadly, selective investigation may identify narrow cases where value is higher.



Single-Boundary or Targeted Deployment Analysis

Further work could test whether smoothing applied only at a specific boundary, or under particular constraint patterns, could deliver a greater marginal impact with lower asset usage.

Full Power-Flow Simulation

Embedding smoothing actions within full network models (e.g., DigSILENT PowerFactory) would allow detailed exploration of how smoothed flows propagate through the network and whether any system states exist where smoothing could unlock additional capacity more effectively.

Asset Behaviour and Real-World Deliverability

The analysis assumes ideal asset response. Should stakeholders wish to explore viability further, a more detailed assessment of deliverability could be considered.

Detailed Modelling of Asset Dynamics

Incorporating latency, ramp-rate limits, power electronics response, and round-trip losses would more accurately represent real-world capability and quantify the gap between theoretical and achievable smoothing.

Locational and Connection-Point Feasibility

A spatial analysis of which assets could influence specific boundaries, taking into account connection voltage, distribution constraints, and Active Network Management schemes — would help assess whether any locality has the right combination of asset availability and boundary behaviour to make smoothing worthwhile.

Market, Commercial and Service Design Development

Given that provider attractiveness is a limiting factor, additional work could explore whether alternative commercial structures could bridge this gap if smoothing were still desired for specific operational reasons.


Stacking with Other Services

A targeted analysis of when and how smoothing might be compatible with frequency response, reserve or wholesale market operation could identify limited circumstances where combined value streams compensate for low standalone revenues.

Data Quality and Scenario Expansion

Improved Constraint Tagging

If smoothing is explored further, improved tagging of constraint-driven BOAs or triangulation via multiple operational datasets would reduce uncertainty in the measurement of potential benefits.




Future System Scenarios

It may be valuable to test smoothing under future network and generation mix scenarios where:


- Wind penetration is higher
- Reinforcement is delayed
- New constraints emerge or existing ones intensify

This may indicate whether the feasibility of smoothing changes under anticipated system trajectories.




7. Glossary

Abbreviation	Description
Adjusted Flow	Power flow with the bids and offers unwound from the metered flow to calculate a flow that precedes any actions taken by the control room
BESS	Battery Energy Storage Systems
BM	Balancing Mechanism
BMU	Balancing Mechanism Unit
BOA	Bid Offer Acceptance
CBA	Cost Benefit Analysis
CCP	Constraints Collaboration Project
Constrained Period	Period where the adjusted flow was above the constraint limit i.e. control room actions had to be taken to manage constraints
Constraint Limit	The assigned power flow capacity of a transmission boundary. If the flow in absence of BOA actions is above the limit, then the boundary is referred to as constrained. Limits are updated on a 10-minute basis
ENCC	Electricity National Control Centre
EWMA	Exponentially Weighted Moving Average
Exceedances	When the boundary flow exceeds the constraint limit
First Difference	The difference between successive values in a time series
GB	Great Britain
Implied Safety Margin	The implied safety margin of power flow below the constraint limit that predicts a certain level of limit exceedance based off the mean flow change distributions.
LDES	Long Duration Electricity Storage
Operational Margin Deployment Cases	The CBA deployment cases that assess the effect of different operational margin for when smoothing benefits are calculated.
Metered Flow	Historical power flow data recorded by meters crossing boundaries



Nested boundaries	A set of transmission constraint boundaries arranged along the same power flow path (e.g. north-to-south), where the same bulk flow crosses each boundary in sequence. As a result, actions at an upstream boundary propagate to downstream boundaries, making their flows and constraints interdependent.
PNA	Power Network Analysis
PV	Photovoltaic
Revised Flow	The metered flow across a boundary with bids unwound to allow additional power through without exceeding the implied safety margin
SMA	Simple Moving Average
Smoothed Flow	Power flow that has been smoothed. For the CBA results the revised flow has been smoothed
Unconstrained Period	Period where the adjusted flow was below the constraint limit i.e. no control room actions had to be taken
VTL	Virtual Transmission Lines: Using storage as a transmission asset to enhance transmission capacity.



8. References

- [1] NESO, “Carbon Intensity Balancing Actions Methodology,” [Online]. Available: https://www.neso.energy/data-portal/carbon-intensity-balancing-actions/carbon_intensity_balancing_actions_methodology.
- [2] BEIS, “Combined Heat and Power – Environmental,” Department for Business, Energy & Industrial Strategy, 2021.
- [3] NESO, “Carbon Intensity API,” [Online]. Available: <https://carbonintensity.org.uk/>.