Phase 1 Demand Forecasting AI-ML

Close out report - Innovation and non-traditional solutions allowance

Commerce Commission

21 September 2026

Contents

1. Introduction

In December of 2025, the Commerce Commission approved Powerco's innovative and non-traditional solutions allowance (INTSA) application to undertake Phase 1 of a multi-phase programme to develop automated, disaggregated electricity demand forecasting capabilities and test whether this capability was appropriate and required to operate as a distribution system operator (DSO). Specifically, the project was focused on leveraging Artificial Intelligence (AI) and Machine Learning (ML) techniques to deliver scalable and automated forecasting aligned with emerging operational and planning needs.

This report provides a summary of Phase 1, how it was delivered, what the results and benefits were, and the key lessons learned. Our hope is that this report can be utilised by other electricity distribution businesses (EDBs), or third parties, looking to explore the use of AI/ML learning in the electricity demand forecasting space. As such, this report satisfies the requirements under Schedule 5.3 of the current default price path (DPP) determination for a close out report. More information on the requirements is provided in Appendix 1.

We are happy to discuss any aspects of this close-out report with the Commission or other interested parties. The first point of contact for this report is Irene Clarke, Policy Manager (Irene.Clarke@powerco.co.nz). Parts of this close-out report are confidential, and we request that these are not published. All confidential information not suitable for release has been marked as such.

2. Original Scope of the Project

On 17 October 2025, Powerco submitted an INTSA application to recover the forecast costs of delivering Phase 1 of a proposed multi-phase Demand Forecasting AI/ML project. The project sought to explore whether AI/ML technologies offer the required robust forecasting capability needed to operate as a DSO and how the Powerco dataset could support this. The project was centred around whether an automated, integrated suite of AI/ML enabled models that can combine large volumes of near-real-time data and generate dynamic insights for short- and long-term planning, is a viable approach to electricity demand forecasting.

Future demand forecasts, particularly when used in a DSO context, carry direct commercial consequences for flexibility service providers and customer investment decisions. Forecasting outputs must therefore be transparent, consistent, defensible and have robust, replicable inputs. They must also capture demand at increasingly granular spatial and temporal levels, whilst reflecting uncertainty in electrification, distributed energy resource uptake and use, wholesale prices, demographic nuances and economic activity. Therefore, network planning and DSO operations require probabilistic forecasts, timely and accurate time-series profiles, and the ability to process diverse datasets at scale, repeatedly and with minimal human intervention.

The project aimed to create a step change in short-term operational and long-term planning forecasts, while reducing manual effort by applying automation and improving demand information. Current deterministic forecasting methods are unlikely to scale successfully to the level required for analysis of very large volumes of data in real time, which draw on a range of different sources and across multiple scenarios. AI/ML offers a way to undertake the processing of data at levels not seen before, and to draw more value from expanding datasets, including weather and socioeconomic information. However, AI/ML's use for electricity demand forecasting remains experimental.

The project is structured into three phases, with this close-out report (and the original INTSA application) covering Phase 1 only:

- Phase 1 – Was designed to test the approach, expose data and modelling challenges, and generate practical learnings on one aspect of short-term forecasting before any decision on attempting to develop a full operational forecasting platform from the ground up. Its near-term value lay primarily in building evidence, capability and market insight of the viability, reliability and practical concerns of directly moving to an AI/ML forecasting approach, rather than evolving a traditional forecasting platform using classical techniques and more manual interventions.

- Phase 2 – Would involve the development and completion of a full short-term forecasting capability, resulting in AI/ML driven and enhanced short-term forecasting across Powerco's network. This would include broader data sources compared to both Powerco's current demand forecasting capability and what was explored in Phase 1, improved model robustness, automated self-validation and outputs integrated into Powerco's operational tools. This phase would test and ensure that an AI/ML demand forecasting platform is successful when integrated into existing business as usual (BAU) tools and deployed as a BAU tool.
- Phase 3 – Would involve the development of long-term forecasting functionality and production integration. It would aim to build a production ready model that included automation, user interface and application programming interface development, multiple scenario support, two-way forecasting (injection and load) and market price elasticity modelling. It would seek to leverage less structured datasets (e.g. information releases on proposed, or recently enacted policy change, discourse on social media) and prove AI/ML capability to effectively inference longer term demand trends from these less direct predictors, especially for sensitivity analysis and strategic scenario planning.

Given that the innovative, experimental, and untested nature of AI/ML demand forecasting tools (and the systems, data, support and other capabilities required to ensure it can be successfully operationalised) is novel in the New Zealand context, a phased project approach was considered most appropriate. This allows the project scope to be adapted as the benefits, risks, and merits uncovered in an earlier phase become clearer. The end of each phase presents an opportunity for reflection and analysis, before a decision is made if the project will move into the next phase.

3. Summary of Phase 1 Delivery

The core work of Phase 1 was carried out over a 12-week period, with a delivery partner selection and project development phase occurring before, and a Powerco review stage occurring after.

The below provides an overview of the stages of the project as they occurred. Other EDBs and third-parties looking to explore similar projects may need to structure such a project in a similar way. The key learnings section of this report (section 7) reflects on the project, how it was delivered, and what Powerco would do differently now that we are privy to the knowledge and experience from having undertaken this work.

3.1 Delivery partner selection

Phase 1 required a combination of electricity-sector knowledge, demand forecasting capabilities, data engineering and science, cloud-platform capabilities and AI/ML expertise and capabilities. Establishing, or enhancing, these capabilities internally for developing a prototype within a time-limited project would have added significant costs to the project, risked scope-creep and significantly increased timeframes. As such Powerco underwent a selection process for a delivery partner and selected Sovereign Systems to deliver Phase 1. Sovereign Systems were able to offer all of the required expertise and capabilities through a partnering model that brought together specialists from Sovereign Systems, AgileData.io, Anthill and Bluestreet Data.

Sovereign Systems provided co-ordination of the specialist capabilities. This was advantageous for an experimental project because it allowed Powerco to test the technical approach and delivery model without committing to permanent internal capability, digital and production infrastructure or ongoing support arrangements.

3.2 Governance and delivery approach

The project involved individuals from across Powerco with DSO, Data Insights, Business Intelligence, Data Engineering, Project Management and Solution Architecture expertise, alongside the Sovereign Systems delivery team.

Powerco provided source data and network and operational context, whilst also sense checking data/outputs. Sovereign Systems co-ordinated delivery and led development of the data-processing, anomaly-detection, profile-generation and prototyping.

Given Phase 1 was intended to test feasibility and expose constraints, this collaborative structure ensured decisions could be informed by both technical evidence and operational knowledge.

3.3 Platform setup and agreed dataset

The first part of the project required the establishment and configuration of the AgileData cloud platform, providing access to the platform for relevant Powerco personnel, and agreement on the datasets that would be used to generate to test the viability of AI/ML demand forecasting outputs.

The agreed datasets included process information (PI) historian measurements, and the network hierarchy and supporting supervisory control and data acquisition (SCADA) alarm and event data. The principal source used for this was Powerco's interpolated 30-minute PI historian data stream from January 2022 to March 2026.

After initial quality filtering, the dataset contained 2,753 measurement tags. These included both direct power measurements and current measurements requiring conversion to megawatts (MW). The source data contained inconsistent unit and voltage naming conventions.

3.4 Data structuring and quality assessment

A material component of delivery was preparing the source data for modelling. A number of discrepancies within the source data needed to be tidied before it could be utilised for modelling. This involved addressing PI historian error states, missing information, identifying genuine zero readings, inconsistent units, inactive circuits, duplicate measurements, sensor failures and gaps in the mapping between measurement points and feeders.

Quality gates were applied before seasonal profiles were produced. A range of controls were applied, including excluding tags where there were fewer than 365 suitable days in the 2024–2025 assessment window, where the measurements were essentially flat or inactive, or where a duplicate measurement of the same circuit was available. MW tags were prioritised over current-derived measurements representing the same physical circuit. These controls ensured that downstream profiles were based on accurate representative data.

3.5 Switching-event and sensor-quality detection

Temporary or permanent network switching can transfer load between feeders and create abnormalities in historian data. This distorts the measurement of actual aggregated and normalised network loading (aggregate customer demand) at the PI network node. If these intervals are not corrected for, it can distort the demand profiles used for forecasting. This was a key challenge in delivering robust forecasting outputs. Unlike the data quality issues explained above, which have occurred due to incorrect input/data capture, the switching event data issues are an outcome of normal network operations

Sovereign Systems implemented five complementary detection methods for identifying abnormal step-down demand, abnormal step-up demand, material departures from a feeder-specific forecasting baseline, corresponding mirrored load step change on connected or nearby feeders, and extended runs of identical readings that indicated a malfunctioning sensor. The methods were applied independently to each feeder but used a common network model to corroborate outcomes through inferencing the equivalent network reconfiguration.

Detection thresholds were expressed as relative measures, including percentiles, z-scores and percentage departures from baseline, rather than fixed MW limits. This allowed the detection to accommodate feeders with different sizes and demand characteristics without the need to manually change parameters.

SCADA alarm and event information was used for corroboration of the network reconfiguration periods detected by the model from within the PI historian data, but SCADA event logs were not sufficiently complete or reliable to serve as the principal data source to detect network reconfiguration directly. The implemented approach therefore demonstrated that switching events could be identified statistically and through general network relationships (e.g. feeders which can be tied) even where a comprehensive historical switching schedule was unavailable. It also confirmed that better event data would improve the treatment of ambiguous cases in a future phase. As an independent activity, we are seeking to improve governance and curation of operational switching logs as this project proved the potential value of this.

3.6 Profile construction

For intervals identified as affected by switching or data-quality issues, replacement values were estimated using predictive models trained on the relevant feeder's historical and non-corrupted information. The corrected dataset was then used to construct seasonal demand profiles for each eligible circuit.

Profiles were produced by season, day of week and 30-minute interval, capturing typical demand pattern and operating ranges. Multiple profiles were produced for each, including mean and upper and lower percentiles. These profile sets provide a statistical view of demand on the network at a granular level, and are essential for future operational, investment and commercial decisions in a DSO context. The outputs were made available through an interactive dashboard in the AgileData platform, enabling Powerco users to view median, minimum and maximum profile information.

In-line with the scope of phase 1 (which only considered part of the short-term forecasting of demand) the delivery produced baseline profiles representative of the current customer demand on the network, rather than long term forward forecasts of future customer needs. Weather normalisation, explicit growth assumptions, smart-meter corroboration and production integration were not part of the Phase 1 implementation. The Phase 1 profiles provide the cleaned and structured demand foundation needed for subsequent medium and long-term forecasting.

4. Results

4.1 Validation, documentation and handover

The outputs were validated against Powerco's 2026 Asset Management Plan feeder forecasts as an independent test of reasonableness. Feeder matching drew on network mapping tables and supplementary feeder-code information. This comparison tested whether the modelled peaks were credible at network scale and identified individual outliers requiring investigation.

The project also delivered implementation documentation describing the technical approach, configuration, source tables, mapping logic, processing methods and known data issues. Powerco received access to the analytical platform and profile outputs.

4.2 Completion against the Phase 1 scope

Sovereign Systems successfully implemented the agreed Phase 1 scope. The platform was established, required datasets were ingested and structured, data quality was assessed, automated switching-event, anomaly and frozen-sensor detection methods were developed, affected intervals were corrected, seasonal time-series profiles were generated, outputs were validated against an independent planning benchmark, and implementation documentation and platform access were provided.

The project achieved its intended purpose as a foundational prototype. It demonstrated that automated data cleaning and demand profiling could be undertaken at Powerco network scale and established the technical evidence needed to inform decisions on subsequent phases. It did not produce a production forecasting service, continuous source-system integration or final business-as-usual operating model, as these were outside the Phase 1 scope.

5. Benefits

Phase 1 produced a functioning, automated data-cleaning and demand-profiling pipeline using Powerco's interpolated PI historian data. Appendix 2 provides examples of these outputs.

The principal benefit is an automated and repeatable way to undertake analysis that would otherwise require manual review of individual measurement points. The prototype improves the quality and transparency of inputs used for forecasting, exposes data and mapping defects, and provides a consistent baseline profile dataset for subsequent short-term forecasting development. It also builds internal understanding of the data, platform and modelling controls needed for AI/ML enabled forecasting, reducing uncertainty before any commitment to production deployment.

Phase 1 is an enabling and exploratory piece of work rather than a complete operational forecasting system. The outputs can support network planning analysis and provide clean historical profiles needed for later forecasting, but further work is required to productionise data flows, resolve data issues, incorporate current-state or recency weighting where demand has permanently changed, use higher-quality switching records to resolve ambiguous events, and add weather normalisation and smart-meter data to produce automated, reliable, granular, operationalised short term forecasts. Long-term forecasts based off this starting point were also outside the Phase 1 scope.

The ability, and underlying capabilities needed to produce reliable, automated profiles of demand, across all seasons and days at full granularity, is a paradigm shift for the industry. This is an essential shift needed to support future DSO capabilities, where there the speed, accuracy and real time nature of forecasts will have significant network and commercial implications given the increased usage of demand flexibility and the higher level of network utilisation (often approaching network limits) targeted under a DSO model. At the same time, distributed energy resource uptake creates greater volatility and uncertainty in that demand. The volume of data required to support this level of temporal and spatial granularity makes manual processing impractical, including for validation and anomaly detection. This reinforces the need for innovative digital approaches to short-term demand forecasting. The proof of concept has delivered considerable value by increasing confidence that an automated capability can be achieved. It has highlighted the key challenges associated with data quality, governance and management, whilst illuminating the need for broader capability uplift across the industry and among prospective vendors and developers.

6. Project Financial Costs Summary

In our application for this work, we had forecast that the completion of Phase 1 would come at a cost of $105,000. Table 1 provides an overview of the actual financial costs associated with this project. The amount in the table is slightly higher than the amount we forecasted. We are providing the table below, and our learnings on financial costs, for the betterment of those looking to undertake similar work.

Broadly, our forecasted costs reflected the amount of work we believed we would need to execute and the cost of that work.

The additional costs of the project, when compared to our forecasts, are due to the amount of work we had to undertake in relation to the tidying of data, and the sharing of data across environments. As per the learnings outlined in this report, although we expected data cleaning to be a part of this work, we had not anticipated that some of this work would require as much investment as it did. This was particularly the case for the sharing of data between, and across environments.

Table 1: Actual costs of the project

Item	Description	Cost
Delivery Partner	The full contract amount for Sovereign Systems to deliver the scope outlined in this report	[redacted]
Internal Resources	The full internal costs of undertaking this work, including expert internal advice, the work of security and data teams, and the cost of writing the close-out report.	[redacted]
Total Costs		$113,000

7. Key learnings and considerations for replication by other EDBs and third parties

There are a number of key learnings that have come out of this work.

The most important learning is around input data. As a result of this project Powerco is continuing its commitment to data governance, curation, management and standards, to ensure adequate data quality to support future business capabilities. A significant number of valuable data sets can already be utilised for demand forecasting inputs, but quality needs a substantial lift in order to achieve the future forecasting requirements of granularity, scale, automation and integration. It was not unexpected that a valuable learning from this project was that data quality was a significant challenge, and the tidying of data took the majority of both Powerco and the Sovereign System team's time and effort.

Part of Phase 1 was to test whether AI/ML tools were already sufficiently developed to the point that, with appropriate training and configuration, they could directly ingest high volumes of untreated or raw data and produce robust, useful forecast outputs without manual intervention. The project clearly evidenced that significant training, configuration and tailoring work is required to ensure raw ML tools can produce adequate insights.

As such, a significant learning from this work is the importance of ensuring datasets are robust. Powerco would recommend that any organisation looking to undertake similar work ensure they have adequate datasets to input into an AI/ML forecasting tool, as well as ensure that those datasets are curated to a high quality. Given the fundamental role of data, a similar project might begin with data discovery and explicit quality gates rather than moving directly to model development.

Powerco is of the view that its delivery approach was strong and is transferable to a range of differing EDB environments. Clear upfront decisions are needed around the nature of what is being developed (is it a proof of concept or is it the development of a production ready tool) and whether the capabilities around that need to sit in-house and be enduring, or if the expertise of an external provider can be used.

8. Further Work and Next Steps

Powerco is undertaking a DSO transformation programme focused on developing the capabilities required to support electrification efficiently and transition to a DSO operating model. Demand forecasting is a foundational capability, and Powerco has therefore explored several forecasting pathways and approaches and solutions in parallel.

The optimal demand forecasting solution, and the shape of the capabilities needed to support Powerco's DSO transformation, remain highly uncertain, particularly given the rapidly evolving digital environment. It was, therefore, expected that some of the exploratory paths and approaches could reach a point in which further investment became less valuable when compared to alternates. Alongside this INTSA funded innovation project, Powerco had continued evolving its own internal forecasting tools and capability and additionally undertook an RFI to understand what commercial off-the-shelf (COTS) software solutions were available.

Powerco has subsequently recently signed a contract for COTS demand forecasting solution. This represents a significant investment in Powerco's demand forecasting capability and will provide a curated and managed forecasting solution through an established platform, with vendor support and scope for ongoing development. The product can be progressively configured and enhanced to support the forecasting needs of Powerco's evolving DSO model. The product already deploys ML functionality, although not to the extent contemplated or tested through this innovation project.

Phase 1 was successful in demonstrating the role that AI/ML models can play in automated data preparation, anomaly detection and demand profiling. However, given the maturity of ML already embedded in COTS solutions, especially for short term forecasting, the value in pursuing with Phase 2 of this project is limited. The insights expected from Phase 2 are likely comparable to those realised through the COTS solution.

Phase 3 was intended to examine innovative applications of AI more to longer-term forecasting, using diverse range of unstructured and qualitative information. Potential sources included Government policy websites, social media discourse, media sites and more. In amalgamation, these sources could generate material insights for long term forecasting such as DER uptake rates, organic growth patterns and potential shifts in customer behaviour. This would move beyond the ML tools deployed for short term forecasting and largely operating with structured numerical datasets and deploying mathematical and statistical algorithms or correlations.

There may, therefore, still be value in exploring this AI approach to long term forecasting in future. Powerco will continue to assess the performance and evolving needs of the COTS, including whether an additional phase of work could use AI/ML to complement or enhance it. At the same time, the AI/ML space is rapidly and constantly evolving. Powerco will be closely monitoring such developments and may decide in future to progress with future phases of this project.

9. Conclusion

Phase 1 of Powerco's AI/ML demand forecasting project achieved its purpose, demonstrating that AI/ML can automate data cleaning, identify network anomalies and produce credible demand profiles at scale. Phase 1 also confirmed that reliable forecasting depends ultimately on high-quality and well-structured datasets, particularly given the core strength of AI/ML models is their ability to process numerous large datasets. The prototype provided Powerco with practical evidence, capability and lessons that will strengthen future forecasting decisions and offer a useful reference for other EDBs considering similar work.

Powerco's newly contracted demand forecasting off-the-shelf product will provide Powerco with the most practical route to an immediate and substantial step change in demand forecasting, backed by an established platform, vendor support and a pathway to evolve forecasting capabilities alongside Powerco's DSO model. Future phases of the bespoke AI/ML programme will therefore not proceed at this stage. Powerco will continue to monitor the performance of our off-the-shelf product, advances in AI/ML and opportunities to incorporate wider datasets, progressing further work only where it can demonstrably enhance the off-the-shelf solution, deliver clear value for customers and the network, or if changes in AI/ML present new opportunities.

Appendix 1 – Check against Schedule 5.3

Table 2: Close out report requirements and how these are meet through this report

Schedule 5.3 requirement for close out report	How the requirement is met
Within 50 working days of delivery of all the INTSA outputs for the project or programme in a non-exempt EDB's INTSA proposal the Commission has approved, the non-exempt EDB must submit a close out report.	This close out report has been submitted within 50 working days of the delivery of the outputs of this work.
(a) whether, and if so how, the project or programme achieved the purpose set out in the INTSA proposal under paragraph (4)(a) and delivered the expected benefits for consumers under paragraph (4)(b);	The key benefits of this work are outlined in sections 5, 7 and 8. Phase 1 successfully established the technical foundations for a scalable, automated demand forecasting capability. As highlighted in this report, and in our application, continued benefits could be achieved through the continuation of this work, however, we do not plan to undertake further work at this stage.
(b) if the project or programme did not achieve the purpose set out in the INTSA proposal under paragraph (4)(a), the non-exempt EDB's view on why the project or programme did not achieve that purpose;	The project achieved the original purpose of Phase 1 outlined in our application, with the learnings informing a decision not to proceed with Phase 2 at this stage. Refer sections 5, 7 and 8.
(c) if the project or programme did not deliver one or more of the expected benefits for consumers set out in the INTSA proposal under paragraph (4)(b), the non-exempt EDB's view on why the project or programme did not deliver the expected benefits;	The project achieved the expected benefits of Phase 1 outlined in our original application. Refer section 5.
(d) any SAIDI INTSA values and SAIFI INTSA values excluded under Schedule 3.1 or 3.2 in relation to the project or programme;	There were no direct SAIDI / SAIFI impacts and no values were excluded related to Phase 1 of this project.
(e) the cause or causes of the interruptions for any SAIDI INTSA values and SAIFI INTSA values excluded under Schedule 3.1 or 3.2 in relation to the project or programme;	N/A
(f) any steps that the non-exempt EDB took to reduce the likelihood or impact on consumers of the interruptions under subparagraph (e); and	N/A
(g) if the non-exempt EDB worked together with 1 or more EDBs to carry out the project or programme, an explanation of how the non-exempt EDB worked together with the other EDB or EDBs to carry out the project or programme; and	Powerco did not work with 1 or more EDBs to carry out this work.
(h) any insights the non-exempt EDB has gained from submitting the INTSA proposal, and from the delivery of the INTSA outputs, that might assist the non-exempt EDB or other EDBs, including with future INSTA proposals.	Refer to sections 4, 5, 6, 7 and 8 of this report.

Appendix 2

The following figures present seasonal half-hourly demand profiles for the same feeder and day of the week across each season. Each profile has been generated from cleansed and normalised historical data and shows the expected demand together with upper and lower percentile bands.

Appendix 2 item 1 — Summer feeder demand profile

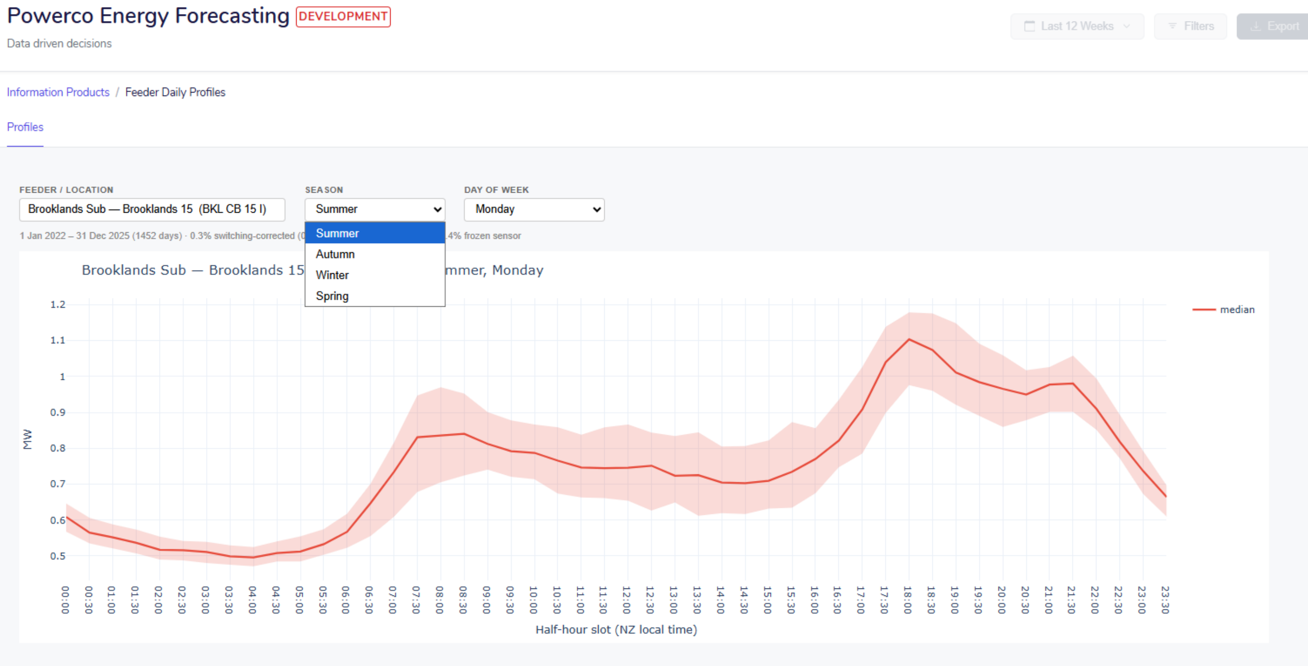

Appendix 2 item 2 — Autumn feeder demand profile

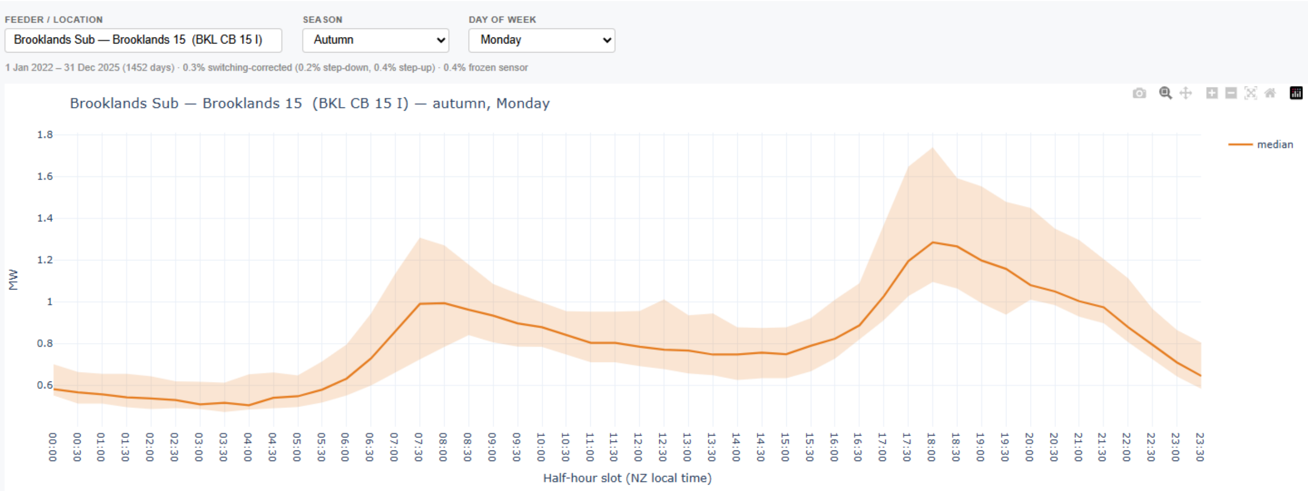

Appendix 2 item 3 — Winter feeder demand profile

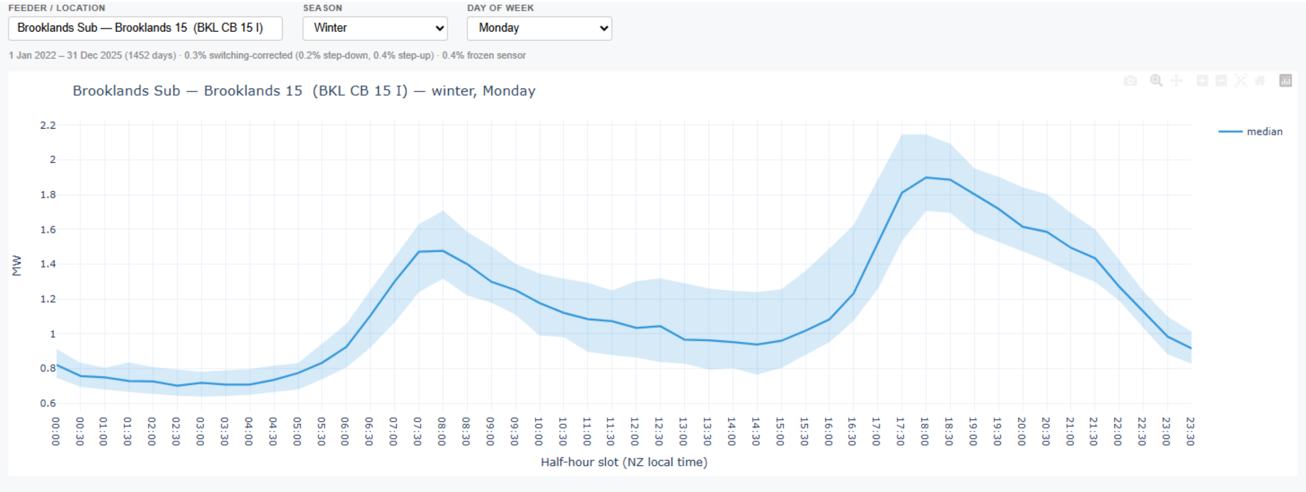

Appendix 2 item 4 — Spring feeder demand profile

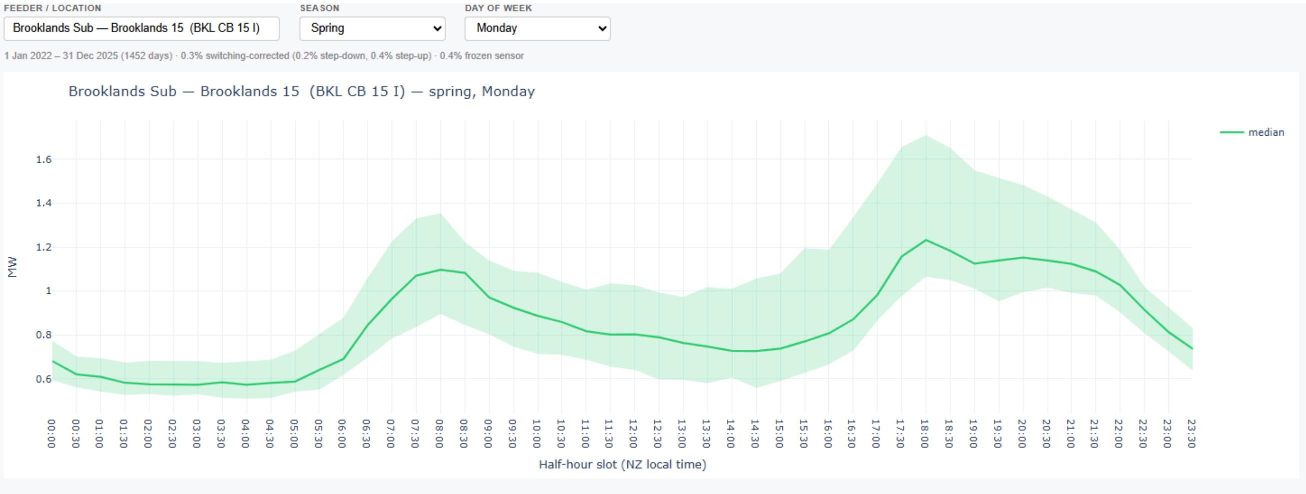

POWERCO