

Estimating the hydro shortfall in a severe dry year and the gap of long-duration firming energy

A report for the Electricity Authority

FINAL REPORT

David Reeve, Stephen Batstone, Corina Comendant
June 26, 2026



Table of contents

Abbreviations	iii
Executive summary	iv
1. Introduction	1
1.1 There is no agreed quantification of the 'dry year problem'	1
1.2 This paper quantifies the dry year problem, and the long-duration firm energy gap	2
2. Determining the hydro shortfall in severe dry periods	3
2.1 The dry year problem is defined by sustained energy shortfalls caused by severe periods of low inflows	3
2.2 Hydro shortfalls reflect worst-case hydro inflows and the resulting storage drawdown	4
2.3 The hydro shortfall in a severe dry period is between 2.4 TWh and 3.4 TWh.....	5
3. Estimating the long-duration firm energy gap.....	6
3.1 Long-duration firm energy reliably responds to sustained hydro shortfalls.....	6
3.2 Determining the need for long-duration firm energy requires an understanding of the supply-demand balance in an average year.....	6
3.3 In an average year, we find that projected demand growth is more than met by current thermal baseload and known generation pipeline	11
3.4 Severe dry periods may create up to a 1.9 TWh long-duration energy gap that the market is unlikely to meet.....	Error! Bookmark not defined.
3.5 The estimated long-duration firming gap could be met through a mix of existing plant, demand response, or other interventions	15
3.6 The winter energy margin standard will not be met at any point in the high-growth high-build scenario.....	16
3.7 Most generation technologies can provide long-duration firm energy, but consistency is what matters	17
4. Principles for designing standardised long-duration firming contracts.....	18
4.1 Long duration firming contracts are needed to underpin investment and manage hydro-related risk.....	18
4.2 Long duration contracts provide clarity and certainty about how firm energy will be deployed	18
4.3 Long-duration firming contracts need to be attractive to buyers, sellers, investors, and analysts	19
4.4 The objectives can potentially be achieved through non-standardised contracts, but standardised contracts may be preferable under some conditions.....	20
4.5 Long-duration firming contracts could lower the risk premium in baseload contracts	20

Appendices

Appendix A	Defining the hydrology deficit.....	21
	Historical hydro shortfall	21
	Adjusting for storage	25
	Capacity and the speed of recovering a shortfall	30
	Commercial and economic.....	32
Appendix B	The relationship between hydro inflows and wind/solar generation	33
	Summary of key points	33
	Description of method	33
	Results.....	35
	Considerations of non-linearity.....	42
Appendix C	Key assumptions underpinning demand growth scenarios	46
Appendix D	Assessment of the winter energy margin	48

Tables

Table 1	– Key characteristics of the hydro shortfall estimations	5
Table 2	– Scenarios of supply-demand balance in an average hydro year.....	10
Table 3	– Scenarios of long-duration firm energy gap	13
Table 4	– The risk of long-duration firming gap given thermal availability.....	15
Table 5	– Interpretation of regression coefficients	34
Table 6	– Key assumptions underpinning demand growth scenarios	46
Table 7	– Electricity demand growth vs 2025 in selected demand growth scenarios	47

Figures

Figure 1	– Annual electricity consumption by scenarios of electricity demand growth	7
Figure 2	– Generation investment pipeline used for this analysis	9
Figure 3	– Annual supply-demand balance in an average year, excluding thermal.....	9
Figure 4	– Demand for thermal in an average hydro year, low-growth low-build scenario	12
Figure 5	– Demand for thermal in an average hydro year, high-growth high-build scenario	12
Figure 6	– Demand for existing thermal in a severe dry year	14
Figure 7	– The long-duration firm energy gap (positive numbers reflect spare thermal)	14
Figure 8	– The WEM in two scenarios of supply-demand balance	17
Figure 9	– Ten worst dry periods, starting in all months, 12-month duration.....	22
Figure 10	– Ten worst dry periods, starting in October, 12-month duration.....	23
Figure 11	– Ten worst dry periods, starting in November, eight-month duration.....	24
Figure 12	– Ten worst dry periods, starting in February, three-month duration	24

Figure 13 – Storage levels from 1980 to 2024, and mean	26
Figure 14 – Cumulative inflows against delta from mean storage and of period starting in February, 8-month duration	27
Figure 15 – Mean storage, minimum and maximum prudent storage.....	29
Figure 16 – Net shortfall (adjusted for storage) for multiple durations ending in September	30
Figure 17 – Capacity to meet shortfall vs duration of deployment and minimum flow limit.....	31
Figure 18 – Worst post-1992 hydro shortfalls	35
Figure 19 – Solar radiation deviation during worst hydro inflows, by start month and duration.....	36
Figure 20 – Wind speed deviation during worst hydro inflows	39
Figure 21 – LOESS: wind shortfall versus hydro shortfall.....	44
Figure 22 – LOESS: solar radiation shortfall versus hydro shortfall	44
Figure 23 – Total annual demand growth by demand growth scenario.....	47

Abbreviations

CCC	Climate Change Commission
DR	Demand response
EDGS	Electricity Demand and Generation Scenarios
GIP	Generation Investment Pipeline
GWh	Gigawatt-hour
MW	Megawatt
SOI	Southern Oscillation Index
SOSA	Security of Supply Assessment
TBA	Projects to be advised
TWh	Terawatt-hour
WEM	Winter energy margin
WV	Water values

Executive summary

The hydro energy shortfall in severe dry years

This analysis finds that the **hydro shortfall in severe dry years can be between 2.4 TWh and 3.4 TWh**, requiring long-duration firm energy to make up for the shortfall during these years. These estimates are based on New Zealand's worst historical inflows and assumptions about hydro storage management:

- **The 2.4 TWh lower bound reflects a minimum-prudence approach to drawing on storage**, managing for the worst dry year sequence while drawing down storage at a rate of 50 per cent of cumulative inflow deficits. This approach draws down all reservoir storage (including contingent storage) in the worst case.
- **The 3.4 TWh upper bound reflects a maximum prudence approach**, managing against the worst inflow sequence of any duration (excluding 1932), and only drawing down normal operating storage. This approach would only draw on contingent storage for the 1932 sequence or a worse sequence than historically recorded.

The long-duration firm energy gap

Thermal generation has historically underpinned baseload supply and provided critical firming during prolonged periods of low inflows, but declining domestic gas supply is eroding this buffer. Although wind and geothermal could potentially contribute to long-duration firming and are progressing through the pipeline, renewables in the current generation pipeline are expected to mainly displace baseload thermal—fully displacing thermal under low demand, and by up to 80 per cent under high demand. In scenarios where thermal is largely displaced, electricity prices will mostly reflect a dominance of renewables, and only infrequently reflect the cost of thermal.

The pace of renewable build-out determines how quickly baseload thermal is displaced, and how much thermal backup remains to cover hydro shortfalls in severe dry years, until renewables can take on a long-duration firming role. As this transition unfolds, **the size of the long-duration firm energy gap depends on how the timing and severity of dry years align with renewable deployment.**

Our estimates of the gap assume a conservative estimate of the total dry year thermal fuel availability. Our results show the following:

- **There is no long-duration firming gap under low demand growth, low generation build and a minimum prudence** approach to storage drawdown for the period we investigated.
- In a high demand growth, high generation build scenario and under maximum prudence reservoir management, **a long-duration firming gap of up to 1.9 TWh** could arise in 2026 and 2027, and **up to 0.5 TWh in 2028 and 2031**, in the most severe dry year. However, we note that at the time of writing, the high demand growth scenario for 2026 appears unlikely. Although the estimated long-duration firming gap could be eliminated through additional renewables, in the near term when the gap is at its largest (2026 to 2027), it will be challenging

to bring additional investment forward. In this case, a package of **contingency measures** could be put in place to largely eliminate this long-duration firming gap if severe dry periods occur.

- Over the medium term, accelerating renewable investment over and above those in the project pipeline would further reduce thermal in a baseload role, liberating thermal plant to move more into a firming role. Over the longer term, further accelerating renewables could significantly play the long-duration firming role.

The winter energy margin (WEM) is aligned with a maximum-prudence framework but **will potentially become increasingly conservative as demand grows. This suggests the WEM should be reviewed.**

The role of contracts market

Closing the long-duration firm energy gap in the high demand scenario requires more investment than the generation investment pipeline is currently showing. New renewables will offset the need for existing thermal to meet baseload needs, increasing its ability to provide long-duration firm energy. Eventually, new renewables could also provide long-duration firming.

Market prices in average hydro years may contain little signal of the need for long-duration firm energy, while investors may reasonably doubt that prices will rise high enough in shortfalls to underpin investment, so **market mechanisms will be required to remunerate long-duration firming plant for availability as well as output.** In particular, contract design should include multi-year terms with regular fixed capacity payments (e.g. option-style fees) to enable capital recovery despite infrequent dispatch, and remain broadly agnostic to technology cost structures. However, the market arrangements around such contracts also need to be carefully designed.

Key recommendations

Our key recommendations for addressing the dry year problem are three-fold:

1. **Ensure that policy affecting investment in new generation and demand response is removing as many barriers as possible,** allowing the generation pipeline to eventuate in a timely manner.
2. **Additional contingency measures to address the long-duration firm energy gap—similar in essence to official conservation campaigns—should remain part of the regulatory toolkit** to manage short-term shortfalls arising from unexpected factors such as delays in new investment, stronger-than-expected demand growth, or more severe inflow conditions than historically observed.
3. **The Authority should commence design of a long-duration firm-energy contract and the surrounding set of design measures,** including the degree to which it should be standardised (if at all), how counterparties to the contract can get equitable access to information in order to value the contract, and how the contracts give the market clarity about how and when firm energy will be deployed.

1. Introduction

1.1 There is no agreed quantification of the ‘dry year problem’

The ‘dry year problem’ is colloquially used to refer to prolonged periods of low inflows into New Zealand’s hydro catchments. Due to the nature of the weather, these periods of low inflows have varying severities in terms of their duration, the amount of hydro ‘fuel’ that the electricity system is deprived of, and the impact on electricity market prices.

Importantly, it is the relatively small size of New Zealand’s hydro storage reservoirs (principally lakes Pukaki, Tekapo and Taupō) that presents challenges for the electricity system in low-inflow periods. Indeed, the dry year terminology is a misnomer; many of the challenging inflow sequences manifest over three to six months, even while the inflows over the year as a whole may be close to average.

As a result, it has proven challenging for the industry to characterise the low inflow periods that cause market stress in simple terms, let alone quantify them. A scan across numerous public documents about the size of the dry year problem reveals numbers ranging from 1.1 TWh to 6.0 TWh, partly reflecting different combinations of inflow shortfalls, access to contingent storage, and assumed or contracted response from some or all current thermal plant. Further, few of these publications clearly explain how the authors used historical inflow and storage data to arrive at the proposed ‘shortfall.’

A good understanding of the hydro shortfall is critical to understanding whether the electricity system has sufficient resources (generation and demand response) that can be activated when low inflows occur for a prolonged period of time. These **long-duration firm energy** resources have a particular characteristic: they can be relied on to produce additional energy in a sustained way over a long period of time; at least weeks, and often months. This service is different from the need to manage short-duration capacity shortfalls—which last only hours and days—which could be provided by batteries, for example.

Today, the long-duration firm energy service is provided by thermal generation—a combination of coal, gas and diesel generation—and some contracted demand response. The availability of fuel for these generators is a critical determinant of their ability to provide long-duration firm energy, as is the extent to which these generators are also being used for another service: baseload generation in an *average* inflow year. This will limit their ability to ‘flex up’ in a dry period.

Although the nature of the hydro problem is unlikely to change in the near term, the mix of generation in the system will change as new investment is commissioned. Depending on demand growth, this changing fuel mix will have a range of implications for how the hydro problem will be managed, not least the fact that new investment may reduce the baseload demands on current thermal generation, improving the system’s ability to respond to low inflows with spare thermal. Equally, fuel may become less available for existing thermal generation, hampering its ability to respond to low inflow periods.

In that context, in pursuit of a secure electricity system we need to ask how we either retain (some) thermal generation in order to continue to provide the long-duration firm energy service, or bring

forward new investment in long-duration firm energy (supply or demand side). Given the infrequent occurrence of dry years, owners or potential investors in long-duration firm energy are likely to need some form of financial contract in order to manage their revenue volatility to an acceptable level. Without this, there is a risk that existing providers exit and/or new forms of firm energy are not brought to market.

1.2 This paper quantifies the dry year problem, and the long-duration firm energy gap

This paper outlines a methodology for how to quantify the size of the problem that prolonged periods of low inflows present to the electricity system, which accounts for both the lower inflows, and how New Zealand's storage reservoirs might be used to mitigate the impact on electricity supply.

We then assess how future demand growth and supply investment scenarios affect the system's ability to respond to severe dry periods, having accounted for the need for current thermal to provide a baseload role in average hydro years, and within a plausible limit on available and flexible fuel. The results of this analysis identify where there is a potential gap between the ability of existing thermal to provide long-duration firm energy, and the needs of a severe dry year.

Finally, we discuss factors that could improve the availability of contracts that could be specifically targeted at helping owners and investors manage the risk associated with long-duration firm energy.

The paper is structured as follows:

- Chapter 2 assesses net energy shortfalls in dry years.
- Chapter 3 estimates the resulting long-duration firming gap—beyond existing thermal and the known generation pipeline under assumed demand growth.
- Chapter 4 sets out the case for new contracts to fill the gap by remunerating long-duration firming plant for availability as well as output.

2. Determining the hydro shortfall in severe dry periods

2.1 The dry year problem is defined by sustained energy shortfalls caused by severe periods of low inflows

The dry year problem is characterised by inflows being consistently below average over varying durations, thus depriving the electricity system of an important source of its fuel. However, the impact of this cumulative inflow shortfall can be moderated by drawing down hydro storage. We call this combined effect on the energy system a **hydro shortfall**.

We recognise that hydro reservoir managers' use of storage in a dry period is influenced by a range of appetites for risk. These conjectural variations need to be taken into account.

We refer to these varying approaches to reservoir management as 'prudence' levels. We have referenced these levels to storage and inflow sequences that end in September. September is chosen because it is, on average, the point at which storage reaches its lowest point over the year, just prior to when increasing spring inflows begin to coincide with falling demand.

We use two prudence levels, minimum and maximum, to describe a less conservative and a more conservative approach to managing storage.

Maximum prudence is defined as managing for the worst historical inflow sequence of any duration (excluding 1932) that ends at the lowest average historical storage level for the year (end of September), while aiming to use only normal operating storage and avoid drawing on contingent storage. Maximum prudence assumes:

- a) on a forward-looking basis, the worst historical inflow sequences remain credible, even when they are preceded by another low-inflow sequence. In effect, this allows for a sequence that is more severe than any observed in the historical record.¹
- b) for each starting month, normal operating storage is depleted to zero under the applicable worst-case sequence.

Further detail on the rationale for this approach is provided in Appendix A.

Minimum prudence is defined as managing for the worst dry year sequence while aiming to fully utilise contingent storage during that sequence. This is achieved by limiting storage depletion to 50 per cent of the cumulative inflow shortfall relative to the historical mean.

¹ Another way of expressing this is to say: if the preceding inflow sequence was the worst sequence in recorded history (around 90 years, so a one-in-90 probability), and yet, looking forward, we select the worst sequence over the coming duration in the 90 years in history (which also has a one-in-90 chance of occurrence), at worst we are simulating a one-in-8,100 inflow sequence.

2.2 Hydro shortfalls reflect worst-case hydro inflows and the resulting storage drawdown

The hydro shortfalls are derived by converting historical hydro inflows and storage into system-wide generation and energy balances, then identifying the worst cumulative inflow shortfalls relative to average across multiple durations and sequences.

The detailed steps involved are as follows:

1. Compile hydro inflow and storage data from Hilltop datasets for the main catchments: Tekapo/Pukaki, Tongariro/Taupō, Hawea/Wānaka/Wakatipu, and Te Anau/Manapōuri.²
2. Model tributary inflows at the point they enter each catchment.
3. Convert inflows into generation (GWh) using linear station-specific conversion factors, cross-checked against WSP reports and JADE input data.³
4. Convert reservoir levels into stored energy (GWh) using linear GWh-per-metre relationships above nominal minimum levels.
5. Use linear conversion factors (GWh per metre) to convert reservoir levels into energy storage, based on metres above nominal minimum levels from Hilltop data.
 - Although some Hilltop storage curves are polynomial, a linear approximation is adopted because it only introduces small errors (generally under 3 per cent, increasing only as storage approaches zero).
 - The relative error increases as storage tends to zero, but we are comfortable that the analysis is suitable for second order approximation of the hydrology deficit. More accuracy would require a substantially more detailed model.
6. Aggregate inflow and storage across catchments to create a system-wide hydrology dataset.
7. Calculate cumulative inflow shortfalls from the mean for durations between two and 12 months, and for different starting months.
8. Select the 10 worst sequences for a range of durations, or for duration and starting month.

More details are provided in Appendix A.

² Hilltop is the hydrological database used for the long-term hydro record. CSV outputs are available from the EA's online datasets.

³ WSP are international engineers responsible for the production, with the associated reports, of the latest inflow and storage data. Jade is the EA's optimisation software that produce water values and water value curves. The conversion factors in its input files were used to check the conversion factors used. The Jade input files are available from the EA's online datasets.

2.3 The hydro shortfall in a severe dry period is between 2.4 TWh and 3.4 TWh

We determine the hydro shortfall to be 2.4 TWh given a minimum prudence level, and 3.4 TWh given a maximum prudence level in a severe dry year.

Table 1 outlines the methodology and key characteristics of the shortfall calculation.

Table 1 – Key characteristics of the hydro shortfall estimations

Prudence level	Historical sequence with the worst net shortfall	Inflow shortfall ⁴	Additional storage depleted in the given inflow sequence ⁵	The hydro shortfall ⁶	Use of contingent ⁷ storage	Maximum baseload-generation equivalent required at 5-6 months in advance	Baseload generation equivalent at 12 months in advance
Minimum	1932	-4.8 TWh over 11 months	2.4 TWh	-2.4 TWh (-4.8+2.4)	Full (~0.5 TWh)	~600 MW over 5-6 months (~2.4 TWh ⁸)	~280 MW (~2.4 TWh)
Maximum	1950	-4.3 TWh over 10 months	0.9 TWh	-3.4 TWh (-4.3+0.9)	Unlikely (but could occur in the worst recorded year)	~700 MW	405 MW

⁴ Cumulative flows below mean inflows over the time period specified.

⁵ Compared to the storage used at that point in time on average. The lowest mean storage (measured on a daily basis) by the end of September would be 1.9 TWh; therefore, in this inflow sequence, the storage level shortfall would be -0.5 TWh, i.e. it has dipped into contingent storage.

⁶ Inflow shortfall adjusted for storage that has been drawn down, according to the prudence level (the minimum prudence uses the characteristic based on the worst inflow sequence, i.e. can use 0.51 of the storage available below mean for every MWh of cumulative shortfall; the max prudence).

⁷ During the course of this project, there has been a draft proposal issued by an expert RMA panel to add contingent storage to normal storage. The final decision is expected by 3 July. For the purposes of this work, we continue to refer to it as contingent storage.

⁸ Waiting any longer than six months to activate any long-duration firm energy means the level of average generation required from that plant to deliver 2.4 TWh will result in the hydro plant being pushed below minimum flows (and therefore spill), at which point the activation of the dry year plant becomes a zero-sum game (i.e. no savings achieved). Hence, dry year plant should not be activated later than five to six months (see Appendix A).

3. Estimating the long-duration firm energy gap

3.1 Long-duration firm energy reliably responds to sustained hydro shortfalls

Fundamentally, the dry year problem requires a fuel and generation source that has the ability to reliably produce 2.4 to 3.4 TWh *more* electricity in a severe dry year than it would produce in an average dry year. Reflecting the extended duration of dry periods, we refer to this as long-duration firm energy.⁹ In the current system, only thermal generation (coal, gas and diesel) and long-duration demand response would provide long-duration firm energy.

Conceptually, renewable generation could provide long-duration firm energy if it had a set of market arrangements whereby: (a) it was only released into the market when hydro inflows deviated below average; and (b) it could, over a long duration, provide a reliable amount of energy. No current forms of renewables (other than storable hydro itself) provide this service today.

3.2 Determining the need for long-duration firm energy requires an understanding of the supply-demand balance in an average year

3.2.1 Summary of steps involved in determining the long-duration firm energy gap

Our analysis must consider whether the existing providers of long-duration firm energy are being used, or will in the future be used, to fill year-to-year mismatches between annual renewable generation and demand in an average hydro year. We refer to this as thermal generation that is providing a 'baseload' generation service. It is from this baseload generation level that the long-duration firm energy providers must flex up by 2.4 TWh to 3.4 TWh in a severe dry year.

In order to assess the ability of the system to deliver this additional electricity in a severe dry year, we need to do the following:

1. Investigate a range of demand scenarios into the future.
2. Consider scenarios of new generation investment.

⁹ Long-duration firm energy does not only apply to extended periods of low inflows. It could apply to extended periods of low wind and low sun. Depending on the future generation mix, it might also be needed to manage the gap between high winter demand and predictably low winter renewable generation (i.e., solar). However, the context of this paper is the hydro 'problem', and we use it in that context.

3. Based on step 1 and 2, determine the quantum of current thermal plant that will be required to meet *annual* demand (i.e. in a baseload role) in an average hydro year, after new generation is accounted for.
4. Determine residual thermal energy available for long-duration firm energy during a prolonged period of low inflows, based on a plausible upper bound on thermal fuel availability less baseload thermal used in step 3.
5. Determine the long-duration firm energy gap, if any, that would result in a severe dry year (the result of step 4, minus the hydro shortfall).


3.2.2 Electricity demand scenarios

In order to present outer ranges of the need for long-duration firm energy, we adopted three demand scenarios:

- For **low demand**, we adopted the Climate Change Commission's (CCC) policy reference scenario.
- For **mid demand**, we adopted Te Kanapu's Aotearoa Electrified scenario.
- For **high demand**, we adopted the Electricity Demand and Generation Scenarios (EDGS) growth scenario.

We acknowledge that these come from different scenarios, and each would have different underlying assumption sets (outlined in Appendix C). However, their use here is simply to stake out plausible outer ranges of demand growth, and include a mid-growth scenario.

Figure 1 – Annual electricity consumption by scenarios of electricity demand growth



3.2.3 The generation investment pipeline

Our analysis of the investment pipeline uses the Electricity Authority's Generation Investment Pipeline (GIP). The GIP distinguishes between three types of generation investments:

- **Committed projects:** A project is committed once a final investment decision has been made.
- **Actively pursued:** If a final investment decision has not been made, but other significant milestones have been reached (e.g. a location being secured in addition to a consent application being submitted or contracts to finance the project executed).
- An "other" category. We did not use these projects.

The GIP projects were adopted for our purposes using the following methodology:


- **Committed projects** are taken at their commissioning dates. Where commissioning dates are in the second half of the year, we defer the project's full annual energy contribution to the following year.
- **Actively pursued projects** were only included if they had an expected commissioning year.¹⁰ We conservatively set the first year of production as the year following their stated commissioning year, to correct for optimism bias.

We only looked at projects through to 2032 because the GIP contained no actively pursued projects beyond this date. That is to be expected, given that the actively pursued status essentially requires either a consent or some form of power purchase agreement. Hence any gap would simply reflect continued demand growth with no investment response from the market, which, for the high demand scenario at least, seems implausible.

The assumed generation pipeline is illustrated in the figure below.


¹⁰ Around 20 per cent of the actively pursued projects (by generation output) had a "TBA" commissioning year. We excluded these projects. Given the requirements for projects to meet the actively pursued definition, we were surprised at the number that met the criteria yet have no estimated commissioning date. We recommend the Authority investigate why this is the case, and make any necessary amendments to the definition of actively pursued.

Figure 2 – Generation investment pipeline used for this analysis



The figure below shows the combination of our assumed supply and demand projections. The figure excludes thermal generation, and we assume that any imbalance would require thermal dispatch (unless new renewable projects over and above those in GIP come online over the period).

Figure 3 – Annual supply-demand balance in an average year, excluding thermal



We observe the following:

- In the near term (2026 to 2027), thermal is still needed under all demand scenarios, as it was in 2025 (~5 TWh of thermal generation).¹¹
- There are just enough committed projects to meet the annual consumption requirements of the low demand scenario over the period 2029 to 2031 without thermal.¹² After this, a small amount of actively pursued projects are required.
- Except in 2030 and 2032, thermal is still required to balance annual supply and demand in the high demand scenario, even with all actively committed projects.

We note that given our aim is to determine the envelope for the long-duration firm energy gap, we excluded the medium-case scenario of demand growth from here on.

3.2.4 Combining demand and supply scenarios

Given that supply and demand projections come from different sources and miss the feedback loop between demand growth and investment, we adjusted the generation pipeline to better align with demand, using our own assumptions. Specifically, we applied a general assumption that market prices—driven by demand—will incentivise new actively pursued generation investment in the following way:

Table 2 – Scenarios of supply-demand balance in an average hydro year

	Demand	Supply
Low-growth low-build scenario	CCC reference scenario	Committed + 10% actively pursued (ex TBAs)
High-growth high-build scenario	EDGS growth scenario	Committed + 90% actively pursued (ex TBAs)

In essence, we are proposing that successive years of low demand are unlikely to incentivise many of the actively pursued projects in the GIP, while successive years of high demand are likely to bring forward most of the projects in the GIP.

As a result, the supply-demand balance in the following charts is somewhat by construction, but will illustrate the reliance on actively pursued (but not committed) projects, and their timing, to maintaining security of supply.

¹¹ This data was sourced from [emi.ea.govt.nz.](http://emi.ea.govt.nz/), which the EA will soon decommission.

¹² This is not the same as saying thermal will not be needed in the system. Some form of firm energy will still likely be needed for peaking, and potentially winter firming.

3.3 In an average year, we find that projected demand growth is more than met by current thermal baseload and known generation pipeline

A key assumption we make here is on thermal fuel availability of 5 TWh, noting that new projects over and above what we have assumed in the pipeline could be brought forward and replace thermal 5 TWh contribution to the system (either as baseload or firming).

To consider how achievable this level of generation is for existing long-duration firm energy sources (i.e., thermal plus long-duration demand response), we need a plausible estimate of what the fuel limitations are for these firm energy sources.

It was beyond our scope to investigate the upper plausible bound on access to coal and gas over the next 10 years. Instead, we observe the following:

- The Reference case in Transpower's draft 2026 Security of Supply Assessment estimated that gas generation in a dry year¹³ would be limited to between 2.0 to 2.5 TWh over the period 2026 to 2035.
- In addition, the Reference case assumed that Huntly would be capable of producing 2.9 TWh over a dry year.
- Taken together, a conservative estimate of the total dry year thermal fuel availability would be 5 TWh.
- We note that, by comparison, total thermal generation in 2025 is also 5 TWh.

Hence, our analysis below compares the resulting requirements for long-duration firm energy 'flex up' in a dry year to a conservative estimate of 5 TWh as an upper bound. Requirements for long-duration firm energy that results in thermal exceeding the 5 TWh upper-bound do not necessarily imply that there will be insufficient fuel; it simply means that fuel arrangements will be a critical factor in achieving that.

The following two charts show that in an average year, projected demand growth is more than met by current thermal baseload and the GIP in both the low-growth low-build and high-growth high-build scenarios.

¹³ The SOSA's period of analysis is the six months from 1 April to 30 September. Our analysis of the hydro shortfall allowed for durations longer than six months, but we adopt these six month generation fuel requirements as they are.

Figure 4 – Demand for thermal in an average hydro year, low-growth low-build scenario





Figure 5 – Demand for thermal in an average hydro year, high-growth high-build scenario



We draw two conclusions from the charts:

- It is likely that, at least after 2028, prices will mostly reflect a dominance of renewables, and only infrequently reflect the cost of thermal. This is important for our discussion regarding contracts below, highlighting that a succession of normal hydro years are likely to result in prices that do not signal the need for firm energy to close the long-duration firm energy gap.

- Despite investment coming to market, the high-growth scenario will require thermal generation over the year that could be between 20 to 75 per cent of actual thermal generation in 2025 (1 TWh to 4 TWh). Demand on thermal at the upper end of this range will leave little ability to increase generation to cope with a dry year, as shown in the next section. This makes understanding the fuel availability for gas plant very important.

3.4 d

The gap is determined as an envelope of low and high estimates, by combining the supply-demand scenarios above with the minimum and maximum prudence approaches to hydro management. We also allow for long-duration firm demand response of 0.5 TWh to occur if a gap is identified.¹⁴


Table 3 – Scenarios of long-duration firm energy gap

	Demand	Supply	Dry year approach
Low long-duration firm energy gap	Low demand growth assumptions; demand response triggered if gap	Committed + 10% actively pursued (ex TBAs)	Minimum prudence
High long-duration firm energy gap	High demand growth assumptions, demand response triggered if gap	Committed + 90% actively pursued (ex TBAs)	Maximum prudence

The chart below shows the demand for existing thermal in a severe dry year compared to a plausible upper bound on its generation potential. The purple bars illustrate the scenario for high long-duration firm energy gap. Any thermal demand above the 5 TWh upper bound is considered the gap, and triggers demand response.


¹⁴ This is based on Meridian's contract with NZ Aluminium Smelters to reduce demand at Tiwai by up to 185 MW for a period of time. The contract allows for four activations of this demand response over a 20-year period; hence we only account for it if a gap arises. We use a conservative proxy for the quantum of energy reduced through this arrangement of 0.5 TWh, reflecting the actual demand reduction observed over the period June 2024 to February 2025.

Figure 6 – Demand for existing thermal in a severe dry year



The chart below illustrates the long-duration firm energy gap, measured as the delta between demand for existing thermal above and the upper bound on thermal fuel of 5 TWh.

Figure 7 – The long-duration firm energy gap (positive numbers reflect spare thermal)



Based on the chart, we observe the following:

- The long-duration firm energy gap is up to 1.9 TWh should a severe dry year occur in the next two years, **and** the system experience a significant growth in electricity demand, **and** reservoirs are managed to the maximum prudence characteristic.

- There is a gap of approximately 0.4 TWh should a severe dry year occur in 2028 or 2031, **and** the system experiences high electricity demand growth between now and then, **and** reservoirs are managed to the maximum prudence characteristic.
- In all other scenarios, current thermal is sufficient to manage a severe dry year, even when using the maximum prudence approach to reservoir management. The quantification of the gap is based on observing the worst hydrological sequence in recorded history. Such a sequence has not been observed since at least the 1970s. Indeed, *all* of the gaps identified would disappear under any inflow sequence except the worst 18 per cent of sequences in the 90 years of inflow history.
- The gaps only occur in the high demand sequence. While we explicitly include high demand because it is a plausible outcome, the gap in 2027 would require a 3 TWh increase in demand in only two years. Such a level of demand increase has not been observed in the 50 years of demand records held by MBIE. This is summarised in the following table.

Table 4 – The risk of long-duration firming gap given thermal availability

	2026	2027-2028	2029 – 2030, 2032	2033-35
Low long-duration firm energy gap	The system is secure only if current thermal plant and required thermal fuel supply is available	There is thermal surplus in the system	There is thermal surplus in the system	There is thermal surplus in the system
High long-duration firm energy gap	Current thermal plant and fuel availability is insufficient to meet maximum prudence	Current thermal plant and fuel availability is insufficient to meet maximum prudence	The system is secure only if current thermal plant and required thermal fuel supply is available	Current thermal plant and fuel availability is insufficient to meet maximum prudence

We note that our assessment above is reliant on:

- new pipeline projects coming to market on the time frames reflect the GIP, and
- maintaining the 5 TWh of thermal, noting that new projects over and above what we have assumed in the pipeline could be brought forward and replace thermal 5 TWh contribution to the system either as baseload or firming, and
- 3.7 TWh of 'TBA' actively pursued projects not being accelerated to the period in the chart.

3.5 The estimated long-duration firming gap could be met through a mix of existing plant, demand response, or other interventions

The gap figures presented above are based on a conservative estimate of fuel availability for thermal plant, rather than a maximum conceivable generation level. Our estimate comprises 3 TWh of coal generation and 2 TWh of gas generation. While gas generation is difficult to predict at the moment,

coal generation could easily be higher than 3 TWh in a prolonged low inflow sequence. If three units at Huntly were to operate for 11 months with 90 per cent availability, they could produce 5.2 TWh alone.

Furthermore, there are a range of contingent measures that if enacted would further reduce the gaps, including:

- use of the 0.5 TWh of contingent storage
- savings arising from an Official Conservation Campaign. These campaigns are enabled in the Code, and if they were ever to be triggered, would surely be triggered by a worst-in-history inflow sequence.

Together with the Security of Supply Assessment's (SOSA) estimate of 2 TWh of gas generation, the extra Huntly generation and contingent storage, the gaps in our worst-ever dry sequence would be eliminated.

Although these responses would significantly alleviate the gap identified by our analysis, the market by itself may not bring all of these outcomes forward.

3.6 The winter energy margin standard will not be met at any point in the high-growth high-build scenario

We have replicated the winter energy margin (WEM) calculation, but replaced generation investment and demand scenarios with those presented above. We have included the WEM assumption that 2 per cent of demand can respond to a dry year scenario (460 GWh).¹⁵

The WEM is consistent with our maximum prudence approach, but is calculated over a six-month period.


Our reproduction of the WEM is broadly consistent with the insights above. However, it does show that our high-growth high-build scenario (including associated investment) would consistently fail to meet the WEM over the entire modelling horizon. This is, in part, due to the fact that the WEM formulation requires that the necessary quantum of long-duration firm energy grows through time proportional to demand.

This has not manifested to date, because demand has not changed since the WEM was developed. We are not aware of any analysis that supports this assertion.

Therefore, our conclusion is that the WEM may be too conservative, effectively implying a level of long-duration firm energy which exceeds the underlying hydro shortfall, even in the maximum prudence case.

¹⁵ This is approximately equivalent to the Tiwai long-duration demand response contract with Meridian.

Figure 8 – The WEM in two scenarios of supply-demand balance



More details on our WEM assessment are provided in Appendix D.

3.7 Most generation technologies can provide long-duration firm energy, but consistency is what matters

Above we have assumed that only current thermal will provide long-duration firm energy during our period of analysis.

In reality, all generation technologies currently used in New Zealand (including geothermal) are capable of being shut down for extended periods (long duration) and only activated in a prolonged dry period. Therefore, they can provide some degree of 'firmness' over extended periods, when used as a reserve form of energy.

However, some forms of intermittent renewables have patterns of generation over the year (e.g., solar) which are anti-correlated with demand. This does not invalidate them as a form of long-duration firming, but when used for periods beyond winter, the implications for hydro storage need to be considered. Solar will produce substantially more in summer than winter.

The selection of long-duration firm energy is therefore an economic problem, not a technological one, that tends to favour lower fixed cost investment propositions (e.g., traditional thermal). However, as the variable costs of these types of firming rise, intermittent renewables may, in fact, become lower overall cost.

4. Principles for designing standardised long-duration firming contracts

4.1 Long duration firming contracts are needed to underpin investment and manage hydro-related risk

In an average year, flexible firm energy will likely be required for short duration peaking, but the prices reflecting this will be limited to peak periods. Further, short duration sources of firming will increasingly meet these peak periods (e.g., direct response and battery energy storage system). These newer forms of short-duration firming are technologically unsuited to long-duration firm energy. As a result, average years are unlikely to signal (through wholesale prices) a need for long-duration storage. If inflow years resembling average hydrology are common, then very high prices are required in a dry period to make a firming investment revenue adequate.

However, when dry years occur:

- prices are often suppressed by policy and political intervention
- prices and settlements are not believed or not trusted by all parties
- infrequent events create different interpretations of risk and volatility.

Even the risk of price suppression in a dry period, where an investment has few opportunities to earn revenue from the wholesale market, will deter investment. Contracts could be struck to encourage investment, however, the same wholesale market problems above mean that contracts cannot be valued nor agreed between parties.

A market design change is therefore needed to underpin the long-duration firming contracts that would both underpin investment and help manage risk.

4.2 Long duration contracts provide clarity and certainty about how firm energy will be deployed

The key thing that long-duration firming contracts need to do is provide clarity and certainty about how firm energy reserves will be deployed in the event of a long-duration hydro shortfall. Thereby, they:

- reduce the inter-year revenue risk associated with firm energy investments, thus improving the investment case
- enable price discovery for dry year firming, driving competition for, and innovation in its provision
- create clearer signals (to hydro operators) about how and when firm energy will be deployed in a dry year, reducing the risk premium in water values

- create clearer signals (to all energy traders) around how firm energy will be deployed, and the certainty of the firm energy reserve, which also reduces premia on other contracts.

4.3 Long-duration firming contracts need to be attractive to buyers, sellers, investors, and analysts

There are some key criteria that must be met for dry year contracts to achieve their purpose:

- The contract must reflect the fundamental characteristics of the dry year problem: infrequent periods of market generation at elevated prices, but still leaving a risk-adjusted capital adequacy shortfall that needs to be recovered through regular fixed payments.
- There must be balanced bargaining power (neither buyer nor seller should have a persistent advantage).
- Counterparties must have equal ability to understand the contract.
- Counterparties must have equal ability to value the contract, including access to trustworthy information about the prevalence of dry years and the nature of potential prices in various dry years.
- The counterparties must be able to confirm that the physical solution underpinning the contract meets the contract needs.
- The contract must be able to sufficiently reduce the investment risk to the risk appetite of the investor. This likely requires:
 - the contract to be long-term in nature (e.g., seven to 10 years)
 - catering for large volumes, i.e. grid scale investment size (from 10 to 200 MW).
- The counterparties must both meet acceptable counterparty credit risk requirements.
- The contract must not unreasonably limit competition for, or innovation in, dry year firming solutions.

Although not a requirement, simpler contract forms that are already used in the industry and are easy to value are likely to have more natural liquidity than contracts that are new and complex. Difficult-to-value contract forms will suffer poor liquidity.

A **swaption** is an example of a contract that reflects the characteristic of long-duration firm energy needs, and is understood by most parties. This will make it easier to value than other more complex forms of firm energy contracts. Further, a physical trigger (knock-in and knock-out) on a standardised swaption will create clarity about how and when the long-duration firm energy will be used. For non-standardised contracts, publication of these triggers could form part of the incentives discussed above. This, in turn, will lower the risk premium in other contract types (explained below).

Although not specifically a contract requirement, we suggest that several cleared counterparties (diversifying counterparty and technology risk) would be preferable.

4.4 The objectives can potentially be achieved through non-standardised contracts, but standardised contracts may be preferable under some conditions

Non-standardised contracts allow a greater deal of innovation and better risk management between parties. To create the clarity and the liquidity necessary to achieve long duration firming, incentives need to be placed on some counterparties to, for example:

- create a distinct shadow price (a firm price on which to base contract valuations)
- create an incentive or obligation to enter contracts
- provide public information on which all parties can rely on for establishing value and risk, including the trigger conditions for incentives (e.g., penalties or obligations)
- correct for imbalances in bargaining power.

Standardised contracts would be considered if incentives could not be designed to achieve clarity, certainty, balanced negotiations, or may still be unlikely to have sufficient liquidity. Standardised contracts would also be considered if published forward prices are a desired output and if contracts are to be made accessible to more parties (e.g. through addressing counterparty credit risk). Although, in this latter case the standardised contract would need to be facilitated through a central clearing process. If the purpose of the standardised contract was to increase liquidity, then the central clearing may also need market making.

Remembering, however, every level of intervention potentially comes at the expense of innovation.

4.5 Long-duration firming contracts could lower the risk premium in baseload contracts

Contracts have emerged in response to demand-driven incentives to contract:

- Long-term prices for demand-driven investments are discovered through:
 - annual baseload contracts for the **general level of electricity demand**
 - seasonal or monthly baseload hedges which signal the price of the **shape of demand over the year** (higher in winter than summer)
 - super-peak contracts which signal the price of the **shape of demand over the day** (higher in morning and evening peak periods than during the rest of the day).
- These contracts, when liquid and understood on both sides, will efficiently signal the cost of meeting demand to the demand side.

As argued above, in the absence of clear price discovery of long-duration firm energy, clear information on how firm energy will be deployed, and the certainty of the firm energy reserve, the risk around price (emanating from the absence of long-duration firm energy) will be included in hydro water values and therefore traders' views on other contracts (principally baseload). As a result, creating clarity on the price of long-duration firm energy, and how it will be triggered, could plausibly reduce the risk premium inherent in other contract prices, especially baseload contracts.

Appendix A Defining the hydrology deficit

Historical hydro shortfall

To define the hydrology deficit in electricity generation, we compiled inflows and storage in GWh by using primarily the Hilltop data available as datasets on the Electricity Authority's website. The major catchments were analysed and aggregated:

- Tekapo/Pukaki
- Tongariro/Taupo
- Hawea/Wanaka/Wakatipu
- Te Anau/Manapouri.

Tributary inflows were modelled as best as possible at the point of entry to each catchment.

Conversion factors were then used based on linear GWh/1000 CMD for all power stations that the respective inflows would flow through.¹⁶ The conversion factors were corroborated with WSP reports and JADE input data.¹⁷


Reservoir storage was modelled on linear conversion factors (GWh/m), based on metres above nominal minimum levels, from the Hilltop data. Some reservoir storage in the Hilltop data is modelled using second and third order polynomials, but our analysis suggested that linear approximation gave little error (generally less than 3 per cent). However, the relative error increases as storage tends to zero. We are comfortable that the linear approximation is appropriate for the modelling done. More accuracy would require a substantially more detailed storage model.

The aggregated inflow data was then organised to look at cumulative deficits from mean for a range of durations (from two to 12 months) for a range of starting months, i.e. the data was organised so that any month could be chosen and the duration (from two to 12 months of cumulative shortfall to mean) for sequences starting on that month could be observed. The data could also compare a chosen duration (from two to 12 months) for any starting month, i.e. every possible two-months, three-months, four-months etc. cumulative shortfall could be observed. The worst 10 any starting month, 12-months duration sequences are shown, i.e. the worst sequences regardless of starting month, in Figure 9.

¹⁶ CMD = cumec.days which is a measure of water volume where 1 CMD = 1 cumecs (cubic metre per second) of water flow for 1 day.

¹⁷ The JADE data and WSP reports can both be accessed on the Electricity Authority's dataset site.

Figure 9 – Ten worst dry periods, starting in all months, 12-month duration




Using every month as a starting month for 93 years of inflow gives over 1,100 separate sequences. The worst sequences shown are less than first percentile outcomes against the whole record. The worst ten range from a 4.25 TWh shortfall to 4.7 TWh, and the most recent sequence this severe is the sequence starting January 1976.

There are questions about the veracity of the 1932 inflow sequence based on known issues with the measurement of river flows from this period. However, as can be seen in Figure 9, the sequences beginning July 1950 and December 1975 are as severe as the sequence beginning February 1932. Sequences from the year 1960 appear more frequently (three times) in the worst 10 than the two 1932 sequences, which ties with 1950 for frequency of occurrence in the worst 10. Overall, 1932 is not an obvious outlier for the annual total, and using July 1950 would give a similar result in any case.

1932 is a difficult year for hydro operators to manage in practice, however. In Figure 9, December 1973 and October 1960 stand out more for fast decline over a short period than February 1932, however 1932 has a long and consistent decline that coincides with the winter period. Therefore, 1932 is excluded from one of the reservoir management characteristics. We explain this further below.

For comparison, we show the 10 worst 12-month duration sequences for sequences that start in October only in Figure 10.

Figure 10 – Ten worst dry periods, starting in October, 12-month duration



We use October as the starting month as we restrict the use of storage to the minimum levels of storage on average, i.e. the lowest expected storage. This storage point occurs at the end of September. The twelve-month October sequences end at the point of minimum expected storage. We discuss this further below.

The selection in Figure 10 reduces the sample size to 93 and the respective percentiles are much higher. The any starting month sequences are dominated by a few long dry periods in the early record. In smaller samples, more recent profiles appear, e.g. the twelve-month shortfalls starting in October feature 2.6 TWh and 2.5 TWh of inflow shortfall respectively for October 2005 and October 2011.

The longer duration inflow shortfalls are dominated by older inflow years, but more recent years feature in shorter durations, for example as shown in Figure 11 and Figure 12.

Figure 11 – Ten worst dry periods, starting in November, eight-month duration





Figure 12 – Ten worst dry periods, starting in February, three-month duration



November 2007 is the worst eight-month duration sequence starting in November, and another three post-2000 sequences also appear in this sample. Six of the 10 worst three-month inflow sequences starting in February are post-2000 years, with the two worst being 2013 and 2014.

Although these events happened over summer and autumn, they strongly shape perceptions of winter risk. They also illustrate how risk aversion can influence decisions; after a 1.9 TWh hydrology shortfall in just three months, it takes a leap of faith to assume inflows will be no more than 2.8 TWh short over the next nine months. The situation looks even more daunting if you are 3.5 TWh short by the end of June 2008 (after a severe eight-month inflow shortfall). In that context, an operator could easily fear another 1.2 TWh deterioration over the following three months, even though the actual outcome was a 0.9 TWh improvement, giving a low but not extreme 2.6 TWh inflow shortfall over the full 12 months.

Adjusting for storage

The inflow sequences demonstrated above are the actual inflow shortages, however, storage can be used to offset these inflows. To assess the range of ways in which storage could be used in a dry period, we applied operating characteristics to reservoir storage. We introduced the concepts of maximum prudence and maximum prudence, which are the most and least conservative approaches of managing reservoir storage that can still be justified as prudent.


We note that, during our analysis, Meridian received draft approval to include contingent storage in its normal operating range. In our analysis we demonstrated historical normal operating storage but sometimes used the ability to use contingent storage, allowing normal storage to go negative on occasion. This terminology will be incorrect if Meridian's storage gets final approval, but the concept of some contingent storage is still useful for our prudent assessments, as we describe below.

We used mean storage as our baseline for expected storage. Mean storage represents the collective decisions of hydro operators who must balance both the risk of inflow shortfalls in dry periods, but also the risk of excessive spill in wet periods. Our prudent approach lead us to use no more than the minimum mean storage expected in a year.

As can be seen in Figure 13, operating levels since 1980 are quite varied. We note that there is no such thing as a normal dry year. Even years that tended to produce average generation in a year will have periods of relatively low inflows and periods of relative high inflows. Also, years that ended up being extremely dry, such as 1932, also had periods of near average inflows. Some dry years also had periods where they were quite wet. The volatile storage sequences in Figure 13 show the effect of operating decisions in response to both actual inflow sequences and perceived risks at the time (of both shortage and spill). The mean storage sequence is an entirely statistical construct; however, we consider that it represents a collective 'neutral' curve. The storage levels at which, collectively, hydro operators have managed around to balance risk of shortfall and risk of spill.

With higher demand and lower winter inflows, there is a reliable trend for storage to generally be highest in the summer and lowest at the end of winter. However, the mean storage level covers a relatively narrow band of storage reflecting the chance that, at any time, hydro operators could be managing periods of above average or below average inflows. Mean storage peaks at the start of February at 3 TWh and reaches a minimum by the end of August which tends to persist over September (at 1.9 TWh).


Figure 13 – Storage levels from 1980 to 2024, and mean



As discussed above, the above chart is normal operating storage. It does not include the 545 GWh of contingent storage at Lake Pukaki.

Actual storage varies for many reasons, but a major driver is whether inflows are higher or lower than normal. To assess the relationship between inflows and storage, we settled on comparing the duration of cumulative inflow sequences (all sequences) against the change of storage against mean at the end of the period. If we look at the period February to September (inclusive), the scatter plot of cumulative inflows and ending delta from mean storage is shown in Figure 14. This period was chosen because it is the period over which storage reliably reduces and gives the strongest correlation between storage and inflow deviation.

Figure 14 – Cumulative inflows against delta from mean storage and of period starting in February, 8-month duration



While the correlation is only 0.37, for most years a linear regression reasonably describes the relationship between cumulative inflows and the delta from ending storage. At least there is not an obviously better function.

It is worth considering whether the worst sequence should be modified for reduced wind and solar output. The analysis in Appendix B finds that solar output has no significant correlation to hydro shortfalls. Generally, the same is true for wind, but wind does show a stronger correlation for three-month periods. It would be prudent to consider an adjustment for wind shortfall for three-month duration hydro shortfalls, but no adjustment for anything longer. As we show below, the prudent characteristics bind for long duration inflow shortfalls, therefore, we have not made this adjustment for the purpose of establishing the net inflow shortfall (net of storage).

To test the limit of prudent operation, we used a linear relationship between cumulative inflows and storage that worked on the basis that the worst recorded 12-month inflow sequence (4.7 TWh) would fully utilise normal operating storage and contingent storage at the lowest mean storage level (1.9 TWh). The coefficient of storage to inflow is then:

$$\frac{1.9 + 0.5}{4.7} = 0.5$$

In analysing the use of storage, it was discovered that an 11-month sequence is actually slightly worse than the 12-month sequence (4.8 TWh), but the 0.5 ratio was still within rounding error and the storage outcome for all sequences was consistent with not exceeding the use of contingent storage.

This may sound like an aggressive use of storage, however, any inflow sequence of 3.7 TWh or less would be catered for using normal storage. The most recent year with inflows worse than 3.7 TWh was

1977. The worst 12-month sequence since 2000 (beginning September 2011) was 3.6 TWh, which would push the bottom of the normal operating range but would only skim into contingent storage, if at all. We considered that this approach would be at the limit of prudence and, therefore, used this approach to define our minimum prudent storage operating characteristic.

As can be seen in Figure 14, a coefficient of 0.5 is far less conservative than the coefficient of 0.28 observed. Higher coefficients are observed for shorter durations, but the reservoir operators seem to put on the brakes on storage decline for longer dry durations. There are a few reasons why this happens, and this operation may be prudently conservative. Some of the reasons are commercial (e.g. portfolio impacts on the hydro operator and the official conservation campaign) but some are physical.

Firstly, if inflows are low and storage has been used and is below mean, then a hydro operator must question whether there is natural mean reversion in climatic processes or whether any historical sequence could occur from that point. While using the worst inflow sequences from that point would imply resulting inflow sequences worse than any previously observed, what would be the basis by which a hydro operator discounted that possibility?

Secondly, which complements the point above, we know that the historical record is likely an incomplete sample of inflow distributions. As recently as 2022, a new worst inflow shortfall was set (worst three-month shortfall beginning in January). And 2007/2008 set new lows for an eight-month duration. Worse inflow shortfalls than historically observed cannot be discounted.

We used the observed conservatism and the two reasons above to guide our definition of the maximum prudent operating storage characteristic.

We have defined the range of net hydro shortfall as being between the range of historically observed inflow sequences and the use of storage defined by the minimum and maximum prudent operating storage characteristics.


The maximum use of storage capacity at the rate of 0.5 of storage (deviation from mean energy stored) per unit shortfall in energy defines our minimum prudent operating characteristic. As discussed above, this is anchored to the end of September, which is when mean storage is at a minimum.

The maximum prudence case is based on assuming that the worst historical sequences are still credible even if the preceding sequence was also low, i.e. implying a worse sequence than historically recorded. The starting storage is defined by setting the storage to the level that under the worst recorded sequence for that duration would use no more storage than normal operating storage. For example, the starting storage for June is the storage equivalent to the worst four-month cumulative inflow shortfall (June to September inclusive). Note, this is the worst any time cumulative shortfall, not just the worst one beginning in June. However, the starting storage is not allowed to exceed mean storage (at which point the risk of spill would also increase above historical operating norms). This only works because the worst 10-month duration sequence is consistent with the characteristics of the longer duration sequences, i.e. holding the storage to mean until the end of December works for all cumulative shortfalls of 10 months or greater, with one exception.

The severe cumulative decline from April to September in the 1932 sequence would require storage to be held at mean for most of the year. It is a judgement call, but we decided that this was too conservative for a prudent characteristic. Leaving one sequence out of the definition of the maximum prudence characteristic is easily catered for as this, or a worse sequence than previously recorded, would be the only sequence that would use contingent storage. Although, as we did not adjust for wind in the sub-three-month sequences, then some use of the contingent storage would also occur in a winter dunkelflaute event for the worst recorded cumulative shortfalls of three months or less.

Using these characteristics gives the storage levels shown in Figure 15.

Figure 15 – Mean storage, minimum and maximum prudent storage



To be clear, these are not necessarily suggested management levels or trigger levels for dry year reserves. Those would depend on the specification of dry period chosen, with the level of conservatism, and both the volume and capacity of dry year reserve available. These characteristics are used to define the use of storage for extreme sequences.

Determining a more specific net dry shortfall, and determining any appropriate physical triggers for contracts, will need the Electricity Authority Board (as the regulator) to determine an appetite for risk. This appetite could sit between the characteristics described above, e.g. such as one in 30 absolute sequences to be managed within normal operating storage.

Determining the appetite for risk is an economic choice (balancing cost of risk of shortfall versus cost of firming and considering price impacts on electricity markets) and is outside the scope of this project. This would be determined through the Security of Supply Assumptions and reflected in the standards.

Figure 16 shows the result of using these characteristics to define the net cumulative inflow shortfall against the worst inflow shortfalls (ignoring 1932 for the maximum prudence characteristic, which would be catered by contingent storage).

Figure 16 – Net shortfall (adjusted for storage) for multiple durations ending in September




Figure 16 shows that the maximum net inflow shortfall (associated with the maximum prudent storage operating characteristic) is defined by the 10-month duration shortfall at **3.4 TWh**. The 11-month duration shortfall defines the maximum net inflow shortfall (associated with the minimum prudent storage operating characteristic) at **2.4 TWh**.

Our assessment of the worst shortfall in hydroelectric generation, adjusted for the use of storage, is 2.4 TWh to 3.4 TWh depending on how conservative the approach to the use of storage.

The final dimension that we considered in defining what might cover for hydroelectric shortfall is the relationship between capacity of hydro-firming reserve versus the duration of deployment.


Capacity and the speed of recovering a shortfall

The recovery of any shortfall is a function of the hydro-firming reserve capacity deployed and the duration of deployment. However, the capacity and duration of hydro-firming reserve also have economic costs and affect optionality. Long durations have little option value. They must be triggered early even though it is still unlikely that the prevailing hydro sequence is an extreme one. This does not necessarily mean that early activation should be avoided if the cost of activation is low (see commercial and economic below). Nevertheless, capacity that needs to be deployed over 12 months is no longer reserve.

There is also a latest point for deployment. As the capacity must increase as the duration of deployment decreases, you can reach the point where the hydro-firming reserve is displacing hydro generation below minimum flow generation. This means that hydro would have to spill to make up minimum flow and no water would be saved.

The trade-off between early activation and the baseload equivalent capacity to make up for net shortfalls of hydro generation up to the limit of minimum inflows is shown in Figure 17.

Figure 17 – Capacity to meet shortfall vs duration of deployment and minimum flow limit



As the lowest cumulative inflows shortfalls of shorter duration are relatively higher than longer duration shortfalls, there is less generation that can be offset before minimum flows are compromised, e.g. the worst two-month shortfall can be severe enough to almost be at the level of minimum flow, whereas over 12 months, even the worst inflow sequences still leave a lot of hydro generation available. Therefore, the capacity limit before minimum flows are compromised is much lower for shorter durations than long durations. In Figure 17, the worst cumulative inflow shortfall is explicitly selected for each duration. This produces the downward sloping trend in Figure 17.

The opposite is true for the prudent reserve capacity. Being able to produce a relatively greater amount of energy in a shorter duration requires increasing capacity.

As the maximum prudence level of reserve is higher at 3.4 TWh, the level of prudent reserve capacity required is higher at any equivalent duration. The latest capacity could be delayed and still recover the hydro shortfall under maximum prudence would be a little before six months before the end of September, at which point around 700 MW of baseload equivalent would be required. This is relatively closely aligned with the current WEM, which is about 3 TWh over April to September at the current level of demand.

Working to the minimum prudence level would delay the need to activate dry period reserve by a little over a month, and then would need around 600 MW of baseload capacity.

Theoretically, delaying an option allows for collecting more information and might avoid some events that turn out to be short duration. However, the value of delaying an option is also a function of both the cost to exercise the option and appetite for risk.

Commercial and economic

Dry period shortfalls, or even the risk of shortfall, are not only physical events—they are also a market price event. The primary impact on prices comes from the need to offset hydro generation. To achieve this, hydro must be priced above the offered price of the dry period reserve for that reserve to get dispatched. Depending on that offer price, market prices may rise significantly.

Prices are also affected by risk and uncertainty. The probability that dry period reserve might be activated can lift Water Values (WV), also leading to higher prices. Uncertainty about dry period reserve around level, offered price, and time of activation can also lift WVs and risk premia.

The appetite for risk can also affect the reservoir operators, with the longer duration lower capacity reserve potentially having lower impact on a portfolio than a higher capacity option, which might push a hydro portfolio to be very high at high prices. However, certainty about the timing, volume and specification of late dry period reserve activation might give confidence about reduced number of activations and might be preferable.

Contract arrangements to procure dry period reserve should also consider the degree to which such arrangements give certainty, balance risk and certainty, and would ideally be based on observable risk factors.

It is also worth considering whether multiple dry period contracts of differing durations might be warranted.

Appendix B The relationship between hydro inflows and wind/solar generation

Summary of key points

Variability of sources of energy:

- The largest hydro shortfalls tend to start in summer or spring and last three months.
- There is no clear pattern for solar radiation deviations by dry year periods (i.e. 12 months).
- There is no clear pattern for wind speed deviations by dry year periods (i.e. 12 months).

Correlation between solar and hydro:

- There is no significant relationship between low hydro inflows and solar radiation, except over dry periods of three months, when low hydro inflows tend to coincide with slightly above average solar radiation (the percentage increase in solar output tends to be ~15 per cent of the percentage of inflow shortfall). However, the relationship disappears entirely for longer durations.

Correlation between wind and hydro:

- On average for durations over three months, there is not a statistically significant relationship between low hydro inflows and wind speed, with Southern Oscillation Index (SOI) being a stronger indicator of wind shortfall (increased shortfall under El Niño).
- The strongest and most significant relationship between reduced inflows and reduced wind output occurs over dry periods of three months, when a statistically significant low-hydro low-wind correlation is observed where the percentage shortfall of wind is ~34 per cent of the percentage of hydro shortfall.

Description of method

Data and time horizon

We have used the following data:

- Hydro inflow data 1932 to 2024 for the four main New Zealand catchments' (Waikato, Waitaki, Clutha, and Manapouri) NIWA weather stations.
- Daily wind speed (m/s) and solar radiation (MJ/m²) from 12 NIWA CliFlo weather stations, 1992 to 2025.
- Monthly SOI data from 1876 to 2022 as an ENSO control:
 - New Zealand's rainfall and wind patterns are shaped by the El Niño Southern Oscillation (ENSO), a naturally occurring climate cycle with two opposite phases. During La Niña, north-easterly winds dominate, bringing wetter conditions to the north and east of the country and drier conditions to the south and west. During El Niño, stronger westerlies

take over, drying out the east while the west receives more rainfall. The same catchments that feed our hydro lakes can be dry under both phases, but for entirely different atmospheric reasons.

- Values persistently above +0.7 indicate La Niña; values persistently below -0.7 indicate El Niño. The further the SOI sits from zero, the stronger the phase. Episodes typically last one to two years and occur on average every three to seven years, meaning any given decade will contain a mix of both phases alongside neutral periods.

Setting up the regression

A dry year period is described by (1) a starting month, which can be any month of the year; and by (2) a duration. We investigated four possible durations: three, six, nine, and 12 months. Combined, this gives a set of 48 possible dry year periods.

For each dry year period, a long-run hydro inflow average was determined using post-1992 data, and a sample was created using the ten years with the greatest inflow deviation from the average.

Similarly, for each dry year period, wind speed and solar radiation shortfalls were calculated as deviations from the long-run average.

SOI was averaged over each period as the ENSO control variable.

Estimated ordinary least squares (OLS) models were used with duration and start month fixed effects, standard errors clustered by year.

$$\begin{aligned} \text{wind radiation shortfall}_i \\ = \alpha + \beta_1 \text{hydro shortfall}_i + \beta_2 \text{SOI}_i + \beta_3 (\text{hydro shortfall}_i \times \text{SOI}_i) + \gamma_d + \gamma_m + \varepsilon_i \end{aligned}$$

Three models for wind and then for solar radiation were run with varying specifications: 1 = no SOI, 2 = SOI included, 3 = SOI interaction included. The results from these were based on the n = 480 of dry year periods identified.

As a robustness check, models were rerun on rank 1 observations (single worst dry period per window combination, n = 48).

Table 5 – Interpretation of regression coefficients

Coefficient	Description
Hydro shortfall coefficient B1	When hydro inflows are X% below their seasonal average, wind speed (or solar radiation) tends to be B1 * X% below its own seasonal average. If hydro is 20% below average and B1 = 0.125, wind is expected to be about 2.5% below average during the same period. This tells us whether a dry hydro period also tends to be a calm or cloudy period.
SOI coefficient B2	For every one unit increase in the SOI (shifting from neutral toward La Niña conditions), wind speed (or solar radiation) tends to be B2 % above its seasonal average, holding hydro conditions constant. When SOI = +1 (mild La Niña) vs SOI = 0 (neutral), we would expect wind to be about 1.22% higher on average. ENSO conditions independently influence wind and solar availability regardless of what hydro is doing.

Coefficient	Description
Interaction coefficient B3	This tells us whether the low hydro-low wind relationship changes depending on ENSO phase. A negative interaction coefficient (-0.135) means the relationship weakens as SOI increases. During La Niña dry years (positive SOI), the link between low hydro and low wind fades, and La Niña tends to maintain or even boost wind availability even when hydro inflows are low. During El Niño dry years (negative SOI), the relationship is stronger, i.e. low hydro and low wind tend to occur together.

Results


Historical deviations of wind, solar and hydro by dry year periods

Hydro inflow deviations

In the chart below, each cell shows the single worst post-1992 hydro inflow shortfall for that start-month and duration, expressed as a percentage below the long-run seasonal average for the same window. The April three-month cell, for example, reports that the driest April to June period since 1992 had inflows 36.3 per cent below normal for that window. By construction all 48 cells are negative: they are the worst shortfall events we observed in the record.

The chart suggests that the hydro shortfalls tend to start in summer or spring and last three months.

Figure 18 – Worst post-1992 hydro shortfalls




Solar radiation deviations

In the chart below, each cell shows how solar radiation deviated from its own seasonal average during the same period as the worst hydro shortfall in that window. For example, the April three-month cell asks: in that April to June period when hydro was 36.3 per cent below average, was solar radiation above or below its own April to June average? Red means solar radiation was below average during the dry hydro period, blue means it was above average, and white means little or no deviation.

The chart shows there is no clear pattern for solar radiation deviations by dry year periods.

Figure 19 – Solar radiation deviation during worst hydro inflows, by start month and duration

Red = below average; blue = above average; year = driest hydro period for that window



Solar radiation full sample

	m_wind_1 Base	m_wind_2 + SOI	m_wind_3 + SOI interaction
Dependent Var.:	wind_shortfall_pct	wind_shortfall_pct	wind_shortfall_pct
hydro_shortfall_pct	0.098 (0.049)	0.125* (0.052)	0.155* (0.058)
avg_soi		1.22 (0.783)	-0.264 (1.27)
hydro_shortfall_pct x avg_soi			-0.135* (0.063)
Fixed-Effects:	-----	-----	-----
duration_f	Yes	Yes	Yes
start_month_f	Yes	Yes	Yes
S.E.: Clustered	by: start_year	by: start_year	by: start_year
Observations	468	468	468
R2	0.16261	0.19511	0.21326
within R2	0.01388	0.05215	0.07352

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1			

- Model 1 (Base): On average, holding duration and start month fixed, there is no statistically significant relationship between hydro shortfall and solar radiation shortfall ($p=0.069$).
- Model 2 (+SOI): On average, holding duration, start month, and ENSO conditions fixed, hydro shortfall remains unrelated to solar radiation shortfall ($p=0.226$), with SOI also falling short of significance ($p=0.072$).
- Model 3 (+SOI interaction): On average, holding duration and start month fixed, neither hydro shortfall nor its interaction with SOI significantly predicts solar radiation shortfall ($p=0.191$ and $p=0.784$ respectively).

Rank 1 solar radiation sample (confined to the driest year in dry year period)

	r1_rad_1 Base	r1_rad_2 + SOI	r1_rad_3 + SOI interaction
Dependent Var.:	rad_shortfall_pct	rad_shortfall_pct	rad_shortfall_pct
hydro_shortfall_pct	-0.257 (0.207)	-0.057 (0.220)	-0.053 (0.216)
avg_soi		3.60** (0.986)	0.924 (1.59)
hydro_shortfall_pct x avg_soi			-0.103* (0.041)
Fixed-Effects:	-----	-----	-----
duration_f	Yes	Yes	Yes
start_month_f	Yes	Yes	Yes
S.E.: Clustered	by: start_year	by: start_year	by: start_year
Observations	48	48	48
R2	0.40141	0.57274	0.58316
within R2	0.12324	0.37419	0.38945

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1			

- Model 1 (Base): On average, holding duration and start month fixed, hydro shortfall does not significantly predict solar radiation shortfall in the most extreme dry year periods ($p=0.219$).

- Model 2 (+SOI): On average, holding duration, start month, and ENSO conditions fixed, a one unit increase in SOI is associated with a 3.60 percentage point increase in solar radiation shortfall ($p=0.002$), with hydro shortfall remaining insignificant ($p=0.797$).
- Model 3 (+SOI interaction): On average, holding duration and start month fixed, the interaction between hydro shortfall and SOI is significant (-0.103 , $p=0.017$), suggesting the solar radiation response to hydro conditions varies with ENSO phase, though hydro shortfall itself remains insignificant ($p=0.808$).

Duration stratified analysis

To investigate whether the hydro-wind and hydro-radiation relationships vary by dry period duration, the models were re-estimated separately for each of the four window lengths (three, six, nine, and 12 months) rather than pooling across durations with fixed effects. This stratified approach allows the slope coefficient on hydro shortfall to vary freely across durations, providing greater within-duration statistical power than the interaction model. Start month fixed effects are retained in each stratified model and standard errors remain clustered by year.

Solar radiation stratified results:

- At the 3three-month window, the hydro shortfall coefficient is -0.147 ($p<0.05$). When hydro is below average over a short acute window, solar radiation tends to be slightly above average, consistent with the anticyclonic blocking conditions that drive short dry spells also reducing cloud cover.
- At six, nine, and 12-month windows, the coefficient turns positive and insignificant, with SOI becoming the dominant predictor at annual timescales (12-month coefficient on SOI: 2.25, $p<0.01$).
- Example: for the three-month stratified radiation model, a 20 per cent hydro shortfall is associated with solar radiation approximately 2.9 percentage points above its seasonal average ($-0.147 \times 20 = -2.94\text{pp}$). This positive relationship disappears entirely at windows of six months or longer. The mechanism driving this pattern is not tested directly in this analysis and warrants further investigation.


Wind speed deviations

In the chart below, each cell shows how wind speed deviated from its own seasonal average during the same period as the worst hydro shortfall in that window. For example, the April three-month cell asks: in that April to June period when hydro was 36.3 per cent below average, was wind speed above or below its own April to June average? Red means wind was below average during the dry hydro period, blue means it was above average, and white means little or no deviation.

The chart shows there is no clear pattern for wind speed deviations by dry year periods.

Figure 20 – Wind speed deviation during worst hydro inflows

Red = below average; blue = above average; year = driest hydro period for that window



Wind full sample

	m_wind_1 Base	m_wind_2 + SOI	m_wind_3 + SOI interaction
Dependent Var.:	wind_shortfall_pct	wind_shortfall_pct	wind_shortfall_pct
hydro_shortfall_pct	0.098 (0.049)	0.125* (0.052)	0.155* (0.058)
avg_soi		1.22 (0.783)	-0.264 (1.27)
hydro_shortfall_pct x avg_soi			-0.135* (0.063)
Fixed-Effects:			
duration_f	Yes	Yes	Yes
start_month_f	Yes	Yes	Yes
S.E.: Clustered	by: start_year	by: start_year	by: start_year
Observations	468	468	468
R2	0.16261	0.19511	0.21326
within R2	0.01388	0.05215	0.07352

signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1			

- Model 1 (Base): On average, holding duration and start month fixed, a one percentage point increase in hydro shortfall is associated with a 0.098 percentage point increase in wind shortfall, though this falls just short of conventional significance ($p=0.056$).
- Model 2 (+SOI): On average, holding duration, start month, and ENSO conditions fixed, a one percentage point increase in hydro shortfall is associated with a 0.125 percentage point increase in wind shortfall ($p=0.024$), with SOI itself not significantly predicting wind shortfall independently ($p=0.131$).
- Model 3 (+SOI interaction): On average, holding duration and start month fixed, the low hydro-low wind relationship is significantly moderated by ENSO phase (interaction: -0.135 , $p=0.040$), with the relationship strengthening under El Niño and weakening under La Niña conditions.

Further model 3 interpretations

Neutral (SOI = 0): $0.155 + (-0.135 \cdot 0) = 0.155$

- On average holding duration and start month constant, during neutral ENSO conditions a 1pp increase in hydro shortfall is associated with a 0.155pp increase in wind shortfall.

SOI = 1 (mild La Nina): $0.155 + (-0.135 \times 1) = 0.02$

- On average, a 1pp increase in hydro shortfall is associated with a 0.02 pp increase in wind shortfall (negligible compound risk).

SOI = -1 (mild El Nino): $0.155 + (-0.135 \times -1) = 0.29$

- On average, during mild El Niño conditions, a 1pp increase in hydro shortfall is associated with a 0.29pp increase in wind shortfall (compound risk).

Rank 1 wind speed sample (confined to the driest year in a dry year period)

	r1_wind_1 Base	r1_wind_2 + SOI	r1_wind_3 + SOI interaction
Dependent Var. :	wind_shortfall_pct	wind_shortfall_pct	wind_shortfall_pct
hydro_shortfall_pct	-0.100 (0.204)	0.257* (0.104)	0.251 (0.119)
avg_soi		6.44*** (1.18)	9.85* (4.09)
hydro_shortfall_pct x avg_soi			0.131 (0.137)
Fixed-Effects:	-----	-----	-----
duration_f	Yes	Yes	Yes
start_month_f	Yes	Yes	Yes
S.E. : clustered	by: start_year	by: start_year	by: start_year
Observations	48	48	48
R2	0.32948	0.68018	0.69097
within R2	0.01201	0.52876	0.54465

signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1			

- Model 1 (Base): On average, holding duration and start month fixed, hydro shortfall does not significantly predict wind shortfall in the most severe dry year periods ($p=0.628$).
- Model 2 (+SOI): On average, holding duration, start month, and ENSO conditions fixed, a one unit increase in SOI is associated with a 6.44 percentage point increase in wind shortfall ($p<0.001$), with hydro shortfall marginally significant ($p=0.021$) indicating ENSO phase dominates wind availability during the most severe droughts.
- Model 3 (+SOI interaction): On average, holding duration and start month fixed, the hydro-wind interaction with SOI is not significant at rank-1 ($p=0.344$), while SOI remains the dominant predictor ($p=0.025$).

Duration stratified analysis

To investigate whether the hydro-wind and hydro-radiation relationships vary by dry period duration, the models were re-estimated separately for each of the four window lengths (three, six, nine, and 12 months) rather than pooling across durations with fixed effects. This stratified approach allows the slope coefficient on hydro shortfall to vary freely across durations, providing greater within-duration statistical power than the interaction model. Start month fixed effects are retained in each stratified model and standard errors remain clustered by year.

Wind stratified results:

- At the three-month window, a one percentage point increase in hydro shortfall is associated with a 0.342 percentage point increase in wind shortfall, significant at the 1 per cent level ($p < 0.001$). This is the strongest and most precisely estimated result in the analysis.
- At six, nine, and 12-month windows, the coefficient declines to 0.125, 0.159, and 0.215 respectively, none approaching statistical significance.
- The same synoptic weather patterns that suppress catchment inflows over three months also tend to suppress wind availability.
- Over longer windows this signal averages out.
- Example: for the three-month stratified wind model, consider a drought where hydro inflows are 20 per cent below their seasonal average. Holding start month fixed, the model predicts wind speed during that same three-month period would be approximately 6.8 percentage points below its own seasonal average ($0.342 \times 20 = 6.8\text{pp}$).

Correlation between hydro inflows and solar/wind

There is no significant relationship between low hydro inflows and solar radiation, except over a three-month period

For a given dry period, we found there is no statistically significant relationship between hydro inflows and solar radiation across most dry year periods and ENSO conditions.

The one exception is over a three-month period where low hydro inflows tend to coincide with slightly above average solar radiation. This effect disappears at longer durations.

There is a correlation between low wind and low hydro inflows during El Niño cycles

Across most dry year period durations there is no meaningful relationship, except over dry periods of three months, where a modest and statistically significant low-hydro low-wind correlation is observed.

Dry hydro periods tend to coincide with below-average wind. The same synoptic weather patterns that suppress catchment inflows over a short window also tend to suppress wind availability.

At longer windows of six months or more, that co-movement averages out.

The hydro-wind relationship is not constant; it varies depending on the ENSO climate phase driving the drought:

- During neutral ENSO conditions, a 20 per cent hydro shortfall over three months is associated with wind speed approximately 3.1 percentage points below its seasonal average.
- Under mild La Niña conditions that effect falls to near zero.
- Under mild El Niño conditions it rises to approximately 5.8 percentage points below average.
- Expanding to the ten driest periods per window brings in many more El Niño.

In years of significantly low hydro inflows, ENSO phase meaningfully moderates wind availability:

- The most extreme dry years have mostly occurred during La Niña or ENSO-neutral cycles.

- Of the 48 rank-1 driest periods across all window combinations, only one falls during an El Niño phase. This does not mean El Niño years are not dry; it means that for any given start month and duration, the single worst drought on record is almost always a La Niña event or ENSO neutral.
- In an La Niña cycle, there is no clear relationship between wind and hydro inflows.
- This means that in extreme dry years, wind is less likely to exacerbate the hydro dry year problem if it is a La Niña event than an El Niño event.
- Less extreme dry years are more likely to be during an El Niño event (than the extreme dry years), when there is a correlation between wind speed and hydro inflows.

The El Niño finding should be treated as directional rather than definitive; the signal is consistent across model specifications but rests on a smaller number of observations than we would like.

Considerations of non-linearity

It is entirely likely that the relationship between hydro shortfall and wind/solar availability (if any) is non-linear in nature. Imposing standard OLS forces a linear relationship when the true relationship may curve. Extreme droughts may trigger fundamentally different atmospheric conditions than moderate ones, producing a disproportionately large effect at the tail. OLS also assumes homoskedastic errors, which may not hold here given that variance in wind and solar shortfall could plausibly increase with drought severity. Standard errors clustered by year partially address this but do not fully resolve it.

Potential assumption violations in using OLS

- **Linearity:** OLS assumes the relationship between hydro shortfall and wind/solar shortfall is linear and additive. There is no a priori reason to believe this. The physical relationship between large-scale atmospheric conditions and renewable resource availability is almost certainly non-linear. A 30 per cent hydro shortfall may not be twice as bad for wind as a 15 per cent shortfall. The data cannot tell us the functional form and imposing one is an untested restriction.
- **Endogeneity:** hydro shortfall and wind/solar shortfall are both outcomes of the same underlying atmospheric conditions. This means hydro shortfall is not strictly exogenous; rather, it is jointly determined with wind and solar by a common driver. SOI partially controls for this but is an imperfect proxy for all relevant atmospheric variation. True exogeneity of the treatment variable cannot be claimed.
- **Heteroskedasticity:** error variance is unlikely to be constant across the distribution of hydro shortfall. Extreme droughts are rarer and likely produce more variable wind/solar outcomes than moderate droughts. Clustering by year addresses some of this but does not fully resolve it, particularly given the small number of clusters (33 years).
- **Overlapping windows/pseudo-replication:** the same calendar year appears up to 12 times across different start months within each duration. This creates within-year correlation in the error term that clustering by year does not fully absorb, since the overlap is within-year not

between-year. This is a form of autocorrelation that standard cluster-robust standard errors do not account for.

- **Small effective sample:** despite 480 observations, the effective independent variation comes from roughly 33 years of data. For the rank-1 sample, there are 48 observations from those same 33 years. This limits statistical power and means results are sensitive to individual influential observations.
- **Treatment heterogeneity:** the effect of hydro shortfall on wind/solar is assumed constant across start months, durations, and ENSO phases in the base model. The stratified and interaction analyses partially relax this but cannot fully account for all dimensions of heterogeneity. A three-month April drought in a La Niña year is a fundamentally different treatment than a nine-month September drought in a neutral year.

Addressing the assumption issues

Standard errors were additionally estimated using a wild cluster bootstrap with 9,999 replications and Webb weights to address the small number of year-clusters ($n=33$). Results were materially unchanged, confirming that the key findings are not an artefact of asymptotic approximations in the clustered standard errors.

The three-month wind result remained the same (0.342 , $p<0.001$) under bootstrapping. The three-month radiation hedge also stayed the same (-0.147 , $p=0.024$).

The pooled wind result actually got slightly more significant under bootstrap ($p=0.013$ vs $p=0.024$ before). The pooled radiation result remains insignificant ($p=0.201$), consistent with all previous models.

SOI in the radiation model is now borderline at $p=0.052$, previously $p=0.072$.


LOESS curve

To test for non-linearity in the relationship between hydro shortfall and wind and solar availability, a LOESS (locally estimated scatterplot smoothing) curve was fitted to the residualised data prior to any parametric estimation. LOESS estimates the relationship between two variables at each point using only nearby observations, weighted by distance, making no assumption about the functional form.

The dashed line shows the OLS linear fit for comparison. Where the LOESS curve deviates from the dashed line, there is evidence of non-linearity that OLS cannot capture. Before fitting the LOESS curve, wind and solar shortfall were each stripped of the variation attributable to ENSO conditions (SOI) and the seasonal and duration fixed effects. What remains is the part of wind and solar shortfall that those factors cannot explain. Plotting this against hydro shortfall isolates the hydro-specific relationship and avoids conflating it with broader ENSO or seasonal patterns that affect all three variables simultaneously.

Wind LOESS


Figure 21 – LOESS: wind shortfall versus hydro shortfall



The wind LOESS plot reveals a meaningful deviation from the OLS linear fit. At moderate hydro shortfalls (around -5 per cent to -20 per cent) the LOESS curve sits close to the OLS line, suggesting the linear approximation is reasonable in this range. At severe shortfalls beyond -20 per cent, the LOESS curve rises above the OLS line but remains below zero. Wind is still below its seasonal average during extreme droughts, but the compound risk effect is weaker than the linear OLS model implies. OLS overstates the severity of compound wind risk at the tail.

Solar radiation LOESS

Figure 22 – LOESS: solar radiation shortfall versus hydro shortfall



The radiation LOESS plot shows the curve tracking close to zero across most of the distribution, with only a mild positive deviation at severe shortfalls (beyond -25 per cent) where radiation tends to be slightly above its seasonal average. The LOESS and OLS lines are nearly indistinguishable across the bulk of the distribution. This is consistent with the regression results and confirms that the null finding for radiation is not an artefact of a mis-specified linear model; the relationship is approximately flat regardless of the functional form assumed.

A further note on non-linearity

To formally test for non-linearity, three additional methods were employed alongside the LOESS: a generalised additive model (GAM) with a penalised regression spline, quantile regression at the 10th through 90th percentiles, and piecewise linear regression.

The LOESS plots provide visual evidence of a potentially important non-linear pattern at the extreme tail of the hydro shortfall distribution. At the most severe drought levels (beyond -20 per cent shortfall), wind remains below its seasonal average but less so than the linear OLS model implies, suggesting OLS overstates compound wind risk at the tail. Solar radiation trends above its seasonal average at these same extreme levels. The GAM produced no statistically meaningful improvement over OLS (wind change in AIC = 0.67, radiation change in AIC = -0.40). Quantile regression revealed the hydro coefficient varies across the wind distribution, but the pattern did not provide cleaner identification than the OLS mean estimate. All formal tests lack statistical power to confirm the tail pattern, only 16 observations fall in the extreme drought category (below -25 per cent shortfall) and confidence intervals at the tail are correspondingly wide. The visual evidence from the LOESS is suggestive but not conclusive.

Appendix C Key assumptions underpinning demand growth scenarios

Table 6 – Key assumptions underpinning demand growth scenarios


Source:	CCC Reference	MBIE EDGS Growth	Transpower Te Kanapu Aotearoa Electrified
In a nutshell	Best understood as a policy-aligned reference scenario: moderate-to-strong electrification, but not an upper-bound grid case.	Best understood as a high economic growth scenario: stronger activity growth across sectors, but not necessarily faster system-wide electrification than more technology-led scenarios.	Best understood as a high-electrification planning scenario: it explores the upper range of grid demand and location-specific load growth rather than a central demand outlook.
Core demand drivers	Demand grows through population, economic activity, EV uptake, building electrification, and process-heat switching under a policy-consistent reference pathway.	Higher GDP growth, higher population growth, increased immigration, increased vehicle kilometres travelled, lower carbon prices, and stronger commercial activity all lift electricity demand.	Strong electrification of end uses is the central driver, with rapid uptake of electric transport, building heat, and process heat under a high-demand electricity future.
Sector emphasis	Transport, buildings, and industrial fuel switching all rise, but generally at a measured pace aligned with emissions-budget settings rather than a “maximum electrification” case.	In the near term, commercial and industrial demand are the main growth drivers; from the late 2030s, transport electrification plays a larger role.	Transport, buildings, and industrial process heat all electrify more deeply, creating stronger total demand and more pronounced regional/grid impacts.
New load assumptions	New loads such as hydrogen and data centres can be represented in ENZ but are included selectively as part of a plausible decarbonisation pathway rather than pushed to the upper end.	Includes explicit datacentre assumptions, with capacity reaching 383 MW by 2029 in Growth, alongside limited hydrogen production assumptions unless stronger switching occurs in other scenarios.	Explicitly tests large and concentrated new loads such as data centres, hydrogen, and other major industrial loads because the scenario is designed for transmission planning.

Table 7 – Electricity demand growth vs 2025 in selected demand growth scenarios

EDGS Growth					
	2030	2035	2040	2045	2050
Commercial	32%	46%	61%	71%	77%
Industrial	11%	26%	39%	44%	45%
Residential	2%	9%	15%	20%	24%
Transport	182%	518%	1096%	1684%	2182%
Transpower Aotearoa Electrified					
	2030	2035	2040	2045	2050
Cooking, heating and cooling	8%	23%	37%	45%	52%
Large industry + data centres	25%	47%	51%	55%	58%
Process heat	59%	282%	501%	569%	591%
Transport	111%	359%	796%	1297%	1758%
CCC Reference					
	2030	2035	2040	2045	2050
Residential	3%	5%	8%	8%	11%
Commercial	1%	2%	2%	2%	3%
Agricultural	2%	3%	4%	4%	5%
Industrial	-6%	-6%	-5%	-5%	-4%
Transport (*)	187%	493%	992%	1529%	1934%

(*) All numbers represent change in electricity consumption versus 2025, except for Transport sector in CCC Reference scenario, which is based on EV proportion of vehicle kilometres travelled.

Figure 23 – Total annual demand growth by demand growth scenario



Appendix D Assessment of the winter energy margin

The winter energy margin is essentially a ratio of:

- the potential generation that could be reliably available over the period 1 April to 30 September, including starting hydro storage (2,750GWh), to
- the expected demand over the period 1 April to 30 September, adjusted for any reliable demand response that would occur in a prolonged low inflow period.

Under the winter energy margin standard, an economically efficient ratio would be between 1.14 and 1.16. In other words, the system should ideally be able to reliably produce 14 to 16 per cent more electricity over the six-month period than would be needed under average hydro conditions.

In replicating the winter energy margin, we adopted all of the WEM assumptions published by Transpower¹⁸ except the following:

- **Electricity demand:** the third-party scenarios we adopted for electricity demand only provided figures for annual demand. We converted these annual figures into April to September demand, using the ratio of 2025 April to September demand to 2025 annual demand. We assumed that this ratio remained the same for our forecasting horizon.
- **Existing generation:** we removed Taranaki Combined Cycle gas generation plant, following its closure by Contact Energy.
- **New generation investment:** we used the same committed investment and proportions of actively pursued investment (by demand scenario) as used elsewhere in this report. We used the SOSA's average capacity factors for wind, solar and geothermal for any new wind, solar and geothermal investment.

¹⁸ At the time of writing, Transpower was consulting on the draft 2026 Security of Supply Assessment. However, this consultation did not include the Final Supplementary Data which is necessary to be able to replicate the underlying assumptions. Hence, we relied on the [Supplementary Data published in the 2025 SOSA](#).

About Sapere

Sapere is one of the largest expert consulting firms in Australasia, and a leader in the provision of independent economic, forensic accounting and public policy services. We provide independent expert testimony, strategic advisory services, data analytics and other advice to Australasia's private sector corporate clients, major law firms, government agencies, and regulatory bodies.

'Sapere' comes from Latin (to be wise) and the phrase 'sapere aude' (dare to be wise). The phrase is associated with German philosopher Immanuel Kant, who promoted the use of reason as a tool of thought; an approach that underpins all Sapere's practice groups.

We build and maintain effective relationships as demonstrated by the volume of repeat work. Many of our experts have held leadership and senior management positions and are experienced in navigating complex relationships in government, industry, and academic settings.

We adopt a collaborative approach to our work and routinely partner with specialist firms in other fields, such as social research, IT design and architecture, and survey design. This enables us to deliver a comprehensive product and to ensure value for money.

For more information, please contact:

David Reeve

Phone: +64 9 909 5810

Mobile: +64 21 597 860

Email: dreeve@thinkSapere.com

Wellington	Auckland	Melbourne	Sydney	Canberra	Perth	Brisbane
Level 9 1 Willeston Street PO Box 587 Wellington 6140	Level 20 151 Queen Street PO Box 2475 Shortland Street Auckland 1140	Level 11, North Tower 80 Collins Street Melbourne VIC 3000	Level 18 135 King Street Sydney NSW 2000	GPO Box 252 Canberra City ACT 2601	PO Box 1210 Booragoon WA 6954	Level 18 324 Queen Street Brisbane QLD 4000
+64 4 915 7590	+64 9 909 5810	+61 3 9005 1454	+61 2 9234 0200	+61 2 6100 6363	+61 8 6186 1410	+61 7 2113 4080

www.thinkSapere.com

independence, integrity and objectivity