Consultation submission form

A proposed reliability obligation to manage dry-year risk

Submission by: Tim Stewart, independent policy analyst and adviser

Submitted to electricitymarkets@mbie.govt.nz by 5pm, 21 July 2026

Submitter information

Name:	Tim Stewart

Privacy of natural persons

Organisation:	Individual

☒ I do not wish my name or other personal information to be included in any information about submissions that MBIE may publish. [Not ticked — I am happy for this submission to be published with attribution.]

☐ I do not want my submission or a summary placed on MBIE's website.

☐ I would like my submission (or identifiable parts) to be kept confidential. [Not ticked — this submission contains no confidential information.]

Questions for the consultation

Q1. Do you agree with the analysis in this chapter of why New Zealand does not currently have sufficient dry-year firming? If not, where do you disagree?

☐ Yes ☒ Yes, with changes ☐ No

I agree with the core analysis: the tenor gap, financing constraints, the investment-timing problem and fuel availability are real, and the observation that consumers pay a dry-year premium year after year even in normal conditions is the most important sentence in the chapter. But the analysis stops short of the structural cause, and the omission matters because it shapes the design choices that follow.

First, the chapter presents under-investment in firming as a market-wide failure of revenue certainty. It is more precisely a failure of incentive alignment under vertical integration. The four vertically integrated gentailers control the large majority of generation and hold a natural hedge against dry years: elevated spot prices cost their retail books but enrich their generation books, so their net exposure to scarcity is small — and for hydro-weighted portfolios, scarcity is net profitable. The parties best placed to finance new firming therefore have the weakest commercial reason to build capacity whose effect is to suppress the scarcity value of their existing assets. Meanwhile the unintegrated participants who genuinely bear dry-year exposure — independent retailers and

independent generators, perhaps a tenth of the market — lack the scale, balance sheet and offtake certainty to finance system-level firming. The tenor gap is binding because the only parties for whom firming is unambiguously value-accretive are the parties least able to fund it. This also explains why the information reforms have not closed the gap: information does not change who profits from shortage.

Second, this structural feature creates a specific risk for the proposed remedy that the problem analysis should acknowledge. Because the obligation falls on the integrated firms and will be met substantially through self-supply, its default effect is to give incumbents a regulatory rationale for holding even more firming capability in-house — while independent retailers, below the threshold, remain wholly dependent on a hedge market whose firm-cover depth the obligated parties control. Unless the design forces cover into tradeable, visible, non-discriminatorily offered form (see Q10, Q11 and Q14), the obligation risks entrenching the very structure that caused the under-investment, with compliance costs passed to all consumers and the benefit of the new firming captured inside incumbent portfolios. Paragraph 35's expectation that the obligation will increase firming access for independent participants is achievable, but only as a design outcome — it is not an automatic property of the mechanism.

Third, the shortfall is framed almost entirely as a deficit of firm generation and fuel. It is equally a deficit of long-duration demand-side capability: New Zealand has essentially one material interruptible load. A deeper pool of contracted curtailable load faces the same hydrology with less unserved risk and no imported fuel. The problem definition should name this, because it shapes the eligibility rules that follow.

Finally, the affordability channel runs through the dry-year premium embedded in forward prices. The success measure for this reform should include an observable decline in that premium, not only improved physical margins (see Q14).

Q2. What do you consider to be the principal remaining barriers to investment in additional dry-year firming following the recent reforms to improve fuel adequacy (Government procurement of an LNG importation facility)?

The LNG facility addresses fuel availability. It does not address revenue certainty, which the paper correctly identifies as the binding constraint. The principal remaining barriers are:

(a) The tenor gap. Nothing in the current framework creates bankable 10–15 year revenue for long-duration firming or long-duration demand response. Lenders will not finance assets whose business case rests on occasional scarcity pricing.

(b) The incentive structure under vertical integration, set out at Q1: the parties best placed to finance firming profit from the scarcity that firming would suppress, while the exposed fringe cannot fund system-scale investment.

(c) No standardised, transferable long-duration cover instrument. The Authority's work on standardised long-duration products should be accelerated, because such an instrument is the natural compliance unit for the obligation and the route to market for independent providers of firming and demand response.

(d) The cost of the marginal cover unit. LNG-backed generation is among the most expensive cover available; unless cheaper cover can compete on equal terms, the compliance price will anchor to imported fuel.

(e) For demand-side firming specifically — the resource class this consultation's own framework contemplates as qualifying cover — the barriers are not primarily financial but regulatory, and they are threefold. Large flexible loads capable of sustained curtailment (flexible industrial processes, grid-scale electrolysis, curtailable data-centre load) require: first, a recognised non-firm connection product — network access priced for load that is contractually curtailable and imposes no peak-driven upgrade costs, rather than the firm-capacity treatment that current connection and pricing frameworks assume; second, access to surplus-period energy on terms that reflect its system value, which depends on contracting structures rather than spot exposure alone; and third, a revenue stream that pays for verified curtailability as a service. None of these exists in New Zealand today, which is why the country has essentially one material interruptible load despite the resource class being proven internationally.

(f) Of those three enablers, this consultation can deliver only the third: a reliability obligation with demand-side eligibility makes curtailability a saleable compliance product for the first time. The first two sit with the Electricity Authority (Code — connection products and network pricing treatment for verified non-firm load) and with Transpower (dispatch protocols and headroom assessment). If they are not progressed in parallel, demand-side eligibility will be formally available but practically unusable — the obligation will recognise a class of cover that the connection framework prevents anyone from building. I recommend MBIE record this interdependency in its advice, and that the updated Government Policy Statement on Electricity — already part of this reform package — signal the expectation that the Authority develop standardised non-firm access arrangements alongside the obligation's implementation. Without that coordination, the eligibility breadth sought at Q10 delivers less than it promises, and the compliance market defaults to fuel and incumbent-held assets.

Q3. Do you agree that, despite the recent reforms to improve fuel adequacy, information, roles and tools, further measures are still needed to strengthen incentives to invest in dry-year firming? If not, please explain why not?

☒ Yes, I agree ☐ I agree in part ☐ No, I don’t agree

Paragraphs 29–30 of the discussion document state the point precisely: better information and clearer accountability do not create bankable revenues for long-duration firming. Information reforms tell the market a shortfall exists; they do not make it anyone's commercial problem to close it. An ex ante obligation converts an identified system risk into contracted demand for cover, which is the missing incentive.

Q4. Do you agree that there is value in building new measures on existing regulatory architecture, rather than creating an entirely new regime from first principles? Why or why not?

☒ Yes, I agree ☐ No, I don’t agree

Yes — for this mechanism. The SOSA, the security standards, the Electricity Risk Curves and the Authority's compliance machinery are established, understood and consulted on, and an

insurance-layer obligation is the kind of measure that existing architecture can carry. A first-principles regime for this purpose would take years and duplicate existing functions.

But the question invites a broader answer, and honesty requires bounding the claim: there are structural features of this market that no amount of fine-tuning within the existing architecture can reach, and the Government should not treat this reform as though it can. The OECD's 2026 Economic Survey of New Zealand describes the structure precisely: market power in this market is episodic, held by two duopolies in succession — hydro-dominant portfolios in wet years (Meridian in the South Island, Mercury in the North Island) and thermal-dominant portfolios in dry years (Genesis and Contact) — with the further finding, drawing on the Commerce Commission's 2025 work, that this control is used to restrict shaped hedge supply to the disadvantage of independents. The Survey's overall conclusion is that New Zealand's electricity prices are structurally too high.

A reliability obligation does not touch that structure. In dry years it may even deepen it: the obligation creates mandated demand for precisely the product the dry-year duopoly sells. The wet-year duopoly — strategic management of hydro offer curves when storage is high — is untouched entirely, since the mechanism only ever activates in the other half of the cycle. The safeguards proposed elsewhere in this submission (demand-side eligibility at Q10, the cover register at Q11, transparency and non-discriminatory offer backstops at Q14) mitigate the dry-year effect; nothing in this package addresses the wet-year effect at all.

New Zealand has recent experience, in another essential-goods duopoly, of accumulating sensible conduct-level measures over successive years while the underlying structure and prices remained essentially unchanged — each measure defensible individually, the set together delivering delay without disruption. This framework should be designed with an explicit guard against repeating that pattern. I therefore recommend a two-track approach. Track one: proceed with the obligation on the existing architecture, as proposed — it is worth doing, quickly, as the insurance layer. Track two: commence design work now on a substantively different market framework capable of addressing the structural duopoly problem the OECD has identified — whether through structural remedies, long-term contracting architectures that separate the financing of new supply from incumbent portfolios, or deeper reform of how firming is owned and priced. The annual reconciliation proposed at Q14 provides the decision gate: if, after a defined period, the explicit cost of the obligation is not accompanied by a measurable fall in the dry-year premium and in the wet-year offer-curve behaviour the OECD describes, the Government should have a developed alternative ready rather than a blank page. Beginning that work in parallel costs little; discovering in 2029 that the fine-tuning track has failed, with nothing prepared behind it, would cost consumers years.

Long-term Obligation

Q5. Do you agree that a long-term obligation should be introduced to strengthen incentives to secure sufficient dependable winter cover in advance where a forecast winter energy adequacy shortfall is identified? If not, why not?

☒ Yes, I agree ☐ No, I don’t agree

Yes. The insurance framing in paragraph 41 is the right one, and converting an implicit, opaque dry-year premium into an explicit, competitively procured cover requirement is a genuine improvement on the status quo.

I would add one honesty point and one condition. The honesty point: an obligation is an insurance mechanism. It should reduce the frequency and severity of scarcity price spikes and it should, if the compliance market is competitive, reduce the risk premium over time. It will not by itself lower the average price level, which depends on the cost of capital and delivery speed of new generation. The framework documents should say this plainly so the reform is judged on what it is designed to deliver. The condition: whether consumers pay less for dry-year insurance overall depends entirely on whether the compliance market is competitive — which turns on the eligibility rules (Q10) and the transparency and concentration safeguards (Q14). Without those, the obligation converts an implicit premium into an explicit cost plus an incumbent margin.

Q6. Do you agree that the long-term obligation should be placed on large wholesale purchasers rather than on, say, generators or the market more broadly? If not, which parties should hold the obligation, and why?

☒ Yes, I agree ☐ No, I don't agree

Yes — purchasers are the right obligation holders. They assemble the portfolio of supply and contractual arrangements on which consumers rely, and critically, a purchaser can discharge the obligation with either supply-side or demand-side cover, whichever is cheaper. A generator-focused obligation (paragraphs 97–105) would bias the response toward supply-side answers and create difficult allocation problems across technologies.

For the reasons set out at Q1, one structural feature needs explicit treatment: the obligated cohort (the four gentailers plus NZAS) contains the parties who are also the dominant potential sellers of qualifying cover, and vertically integrated purchasers will meet much of their obligation through self-supply. Two consequences follow. First, self-supplied cover must face exactly the same evidential standard as third-party cover — no presumption of firmness because the asset and the obligation sit in the same corporate group. Second, the Authority should publish the volumes and volume-weighted prices of arm's-length cover transactions each cycle, so the independent segment of the compliance market is visible and priced, and the cost the integrated firms attribute to self-supplied cover can be benchmarked against it. Third, because the obligated purchasers are also the only material buyers of new cover, the buyer-side safeguards proposed at Q14(7) are integral to the purchaser-based design rather than optional refinements.

Q7. Do you agree that a de minimis threshold should apply to limit the obligation to wholesale purchasers with a material winter load? If so, is the 500 GWh threshold appropriate, and should it be measured by reference to aggregate load over the winter period?

☐ Yes, I agree ☒ I agree in part ☐ No, I don't agree

A threshold is appropriate, winter-period aggregate load is the right measure, and 500 GWh is a reasonable starting point. Three refinements: (a) the anti-fragmentation rule signalled at paragraph 61 should aggregate related-party and common-control load, or the threshold invites structuring; (b) parties newly crossing the threshold should get a defined phase-in (e.g. first obligation cycle at

reduced allocation) so the threshold is not a growth penalty on expanding independent retailers; (c) the threshold should be reviewed on a fixed cycle (say five-yearly, aligned with the security-standard reviews) as retail market shares evolve.

Q8. Do you agree with the proposed staged approach under which a warning notification would be issued five years out, a binding obligation would be issued three years out and compliance would be assessed one year out? If not, what timing or trigger approach would be preferable?

☐ Yes, I agree ☒ I agree in part ☐ No, I don't agree

The staged structure is sound. My concern is the gap between the three-year binding horizon and the delivery time of most new long-duration resources: new firming plant, new long-duration storage or a new large flexible load is typically a four-to-six year proposition through consenting, financing, equipment procurement and construction. If the binding obligation crystallises only three years out, the compliance response will be dominated by whatever already exists — largely incumbent-held assets and fuel — rather than by new investment, which defeats the investment purpose described at paragraph 38.

Two adjustments would fix this. (a) Allow the Authority to issue the binding obligation up to five years out where the forecast shortfall is large or persistent across successive SOSA assessments. (b) Provide safe-harbour crediting for early action: from the warning notification, the Authority should operate a pre-registration process under which a proposed cover arrangement can be certified — as to qualifying status and crediting value — on the rules as they stand, with that certification binding at the later obligation and assessment stages even if eligibility definitions or weighting methodologies subsequently evolve, and with certified cover transferable between obligated parties. Without this, a purchaser acting at the warning bears not only the commercial risk that the obligation never crystallises (which it should bear) but also the regulatory risk that its cover is later classified or credited differently (which it should not); the rational response is to wait until the binding notification and comply with existing assets. That defeats the investment purpose of the regime, because the warning stage is the only point with enough lead time for new resources — and it denies new-build developers the certified, financeable compliance product that would let them raise capital against the obligation. The Great Britain Capacity Market's prequalification process is the working precedent.

Q9. Do you agree that any identified adequacy shortfall should be allocated among obligated wholesale purchasers by reference to their shares of expected winter demand? If not, what allocation approach would be preferable?

☐ Yes, I agree ☒ I agree in part ☐ No, I don't agree ☐ Not sure/no preference

Allocation by expected winter demand share is the right basis. Two design points: (a) the grossing-up to 100 per cent of the shortfall (paragraph 76) means obligated parties carry cover attributable to sub-threshold load — acceptable on proportionality grounds, but the resulting cross-subsidy should be quantified and disclosed, since it will be recovered from the obligated parties' customers; (b) customer churn between allocation and assessment needs a mechanism. The cleanest answer is to make compliance positions transferable: if obligations attach to standardised, tradeable cover instruments (see Q10 and Q14), a purchaser losing load can sell surplus cover to the purchaser gaining

it. Transferability solves the switching problem and simultaneously creates the secondary market that makes cover prices observable.

Q10. What forms of cover should be eligible to meet the long-term obligation?

This is the single most consequential design choice in the regime, and I urge MBIE to settle it in the purchaser's favour of breadth. Eligibility should be capability-tested and technology-neutral: any arrangement that verifiably delivers energy to the system, or verifiably removes demand from it, for the relevant duration and at the relevant notice, should qualify. Paragraph 81 already points this way ("generation, demand response, contractual rights or other arrangements"); that breadth should be locked into the legislation and Code, not left to be contested instrument by instrument.

Eligible cover should expressly include: existing or committed firm generation with demonstrated fuel arrangements; long-duration storage; contracted demand response; contracted interruptible or curtailable load, including large flexible industrial loads capable of sustained multi-week curtailment; batteries, credited strictly within their duration capability; and standardised transferable long-duration cover contracts.

On interruptible load specifically: from the system's perspective in a dry winter, a verified megawatt-hour of load reduction is indistinguishable from a megawatt-hour of firm generation — and it requires no imported fuel. New Zealand already relies on precisely this logic through the Tiwai demand-response arrangements, which the discussion document's own framework treats as material cover. The regime should generalise that logic. Internationally, reliability mechanisms that began generation-only have converged on demand-side eligibility (the Australian RRO's contracting arrangements and the RERT, the GB Capacity Market, PJM and other North American capacity markets). A new class of large flexible loads capable of sustained curtailment — including grid-scale electrolysis and curtailable data-centre load — is emerging internationally; an eligibility definition that excludes it would draft the regime around the previous decade's technology set at exactly the moment new long-duration demand-side cover is becoming available.

Two disciplines should accompany the breadth. First, duration-weighted crediting: the document identifies the problem as up to 3 TWh over a three-month dry period — an energy problem, not a peak-capacity problem — so a resource able to contribute for twelve or more weeks should earn materially higher credit per megawatt than one exhausted in hours or days. This one feature does most of the work of ensuring the obligation procures the cover the system actually needs, and it correctly positions batteries without excluding them. Second, additionality and unique attribution: qualifying cover must be additional to, or verifiably identified within, the SOSA Expected Scenario, and each unit of cover must be allocated to exactly one obligated party per period (see Q11 on the registry).

Q11. How should the compliance framework be designed? In particular, what evidential standard should apply, and how should the regime address more difficult issues such as hydro treatment, attribution, portfolio treatment, shared arrangements and double counting?

Evidential standard: "credible deliverability", assessed on three limbs — physical (plant plus demonstrated fuel, storage or water), contractual (a clear, enforceable entitlement with firmness warranties and, ideally, an indemnity covering penalty exposure on counterparty non-performance),

and operational (demand response and interruptible load must be readiness-tested, not merely documented — a periodic activation test, on the model of AEMO's RERT readiness verification, should be a condition of continued crediting).

Hydro: the Authority should set conservative deemed-firmness values at catchment level, using hydrological assumptions consistent with the conditions the regime is designed to insure against — not the obligated party's own optimism. The asymmetry principle in the short-term chapter (paragraph 152) is right and applies equally here: if thermal must demonstrate fuel, hydro must demonstrate water on equivalently sceptical assumptions.

Attribution, shared arrangements and double counting: these are all solved by the same instrument — a central register of qualifying cover, held by the Authority, in which each unit of cover for each compliance period is uniquely allocated to one obligated party. Shared arrangements are then handled by explicit registered splits rather than by after-the-fact argument. The register also carries the seller-side consequence identified at paragraph 161: a generator that sells cover to one party thereby increases the load its own portfolio must stand behind, and the register makes that visible automatically. Portfolio treatment: diversity credit for geographically dispersed intermittent portfolios is legitimate but should be capped at conservative capacity-factor values set by the Authority, not self-assessed.

Q12. Should the regime deal with arrangements that appear compliant in advance but later fail because of counterparty failure, project delay, force majeure or similar events? If so, how?

☒ Yes, I agree ☐ I agree in part ☐ No, I don't agree ☐ Not sure/no preference

Yes. The right structure is a re-procurement duty rather than an automatic penalty: on a notified failure, the obligated party enters a defined cure window within which it must replace the cover, with penalties attaching to failure to re-procure (or to misrepresentation at the original compliance date), not to the supervening event itself. Force majeure should be drawn narrowly, distinguishing genuinely exceptional events from ordinary commercial risks that prudent contracting should manage — the document's caution at paragraph 95 is well placed. Two supporting features: material cover contracts should be required to include performance warranties and penalty indemnities (which puts the performance incentive on the party actually controlling delivery), and the cover register should carry status flags so that at-risk cover is visible to the Authority early rather than discovered at assessment.

Q13. Should penalties for non-compliance be designed and calibrated to tilt investments to comply without overly encouraging inefficient over-procurement? Please explain your reasoning.

☒ Yes, I agree ☐ I agree in part ☐ No, I don't agree ☐ Not sure/no preference

Yes — the "tilt" calibration in Table 3 is the right target, and the dollar-per-MWh shortfall charge at paragraph 93 is the right instrument. The rate should be set modestly above the cost of the best available qualifying cover and published in advance. This has a second, under-appreciated function: a published shortfall rate acts as a de facto price cap on the compliance market. No rational purchaser pays more for cover than the penalty for not holding it, so the penalty rate protects consumers against monopoly pricing of cover in a concentrated market — which, given the small number of potential sellers, is a real risk. Escalation for repeated shortfalls is appropriate. Penalty revenue should be

recycled to consumers or applied to backstop cover procurement, not retained as general revenue, so the regime cannot drift toward being revenue-positive for the Crown.

Q14. What other design features, transitional arrangements, or implementation issues should be considered to ensure that the long-term obligation is workable, proportionate, cost-effective, and beneficial for consumers?

Five features would materially strengthen the regime:

(1) An annual public reconciliation. Each year the Authority should publish the explicit cost of the obligation alongside its estimate of the dry-year risk premium embedded in forward contract prices. A well-designed obligation should show that premium falling as credible cover accumulates. If the explicit cost rises while the embedded premium does not fall, the mechanism is adding cost rather than converting it — and that signal should be visible early. This single metric is the honest public test of whether consumers are benefiting.

(2) Concentration safeguards in the compliance market. The natural sellers of cover are the incumbent integrated firms whose assets benefit most from scarcity. The Authority should monitor concentration in cover provision, publish volume-weighted cover prices each cycle, and hold a backstop power to require holders of demonstrable surplus qualifying cover to offer it on published, non-discriminatory terms where the market shows withholding.

(3) Standardised, transferable cover instruments, aligned with the Authority's existing work on standardised long-duration products. Standardisation makes compliance auditable, transferability solves customer churn (Q9), and together they create the entry route through which independent providers of firming and demand response can sell into the compliance market.

(4) Discipline in the LNG interaction. The obligation should not be calibrated — explicitly or in effect — to guarantee terminal throughput. LNG-backed cover is a legitimate deep backstop; duration-weighted, technology-neutral eligibility ensures cheaper cover clears first and imported fuel sets the ceiling rather than the clearing price.

(5) Settle demand-side eligibility in the primary legislation. The amendments to the Electricity Industry Act 2010 that establish the framework should state expressly that demand response and contracted interruptible load are within the scope of qualifying cover, so the point is settled in the Act rather than re-litigated in Code consultations instrument by instrument.

(6) Tenor for new resources. The obligation as designed generates winter-by-winter compliance demand, but the resources capable of closing a structural shortfall — new firming plant, long-duration storage, high-capex flexible load — require revenue tenors of a decade or more, as the discussion document itself acknowledges at paragraph 100. This mismatch is the documented weakness of Australia's RRO, which has tended to reshuffle existing contracts rather than finance new capacity. Great Britain's Capacity Market resolved it by offering new-build resources agreements of up to 15 years while existing resources receive one year. The New Zealand analogue: where the SOSA shows a shortfall persisting across successive assessments, the Authority should be empowered to certify multi-winter cover — a single arrangement, certified once as to qualifying status and crediting, covering a defined strip of future compliance periods, transferable between obligated parties. Without a tenor mechanism, the compliance market will clear against existing incumbent-held assets

by default, and the regime will procure insurance without ever financing the new resources that would reduce the premium it is insuring against.

(7) Buyer-side concentration. The safeguards above address the risk that incumbents withhold or overprice cover as sellers. The obligation also creates a buyer-side risk that deserves explicit treatment: it makes a small number of vertically integrated purchasers the mandated buyers of essentially all qualifying cover, so every new independent firming or flexibility resource faces a handful of possible customers — each of which is also its competitor in the flexibility market. Left undesigned, the likely equilibrium is that monopsony bargaining suppresses the margins of independent developers, the cover that is built is contracted into incumbent portfolios where its dispatch and hedge-market offering are controlled by the incumbent's book, and the regime ends by organising the nation's new dry-year response under the control of the same firms whose market position it was meant to discipline — with compliance itself supplied as the justification.

Three design features would cut this loop. First, unbundle the compliance attribute from operational control: cover should be transacted as a standardised registered instrument under which the obligated buyer receives the compliance credit while the provider retains operational control of the resource, dispatched under its registered protocols — meeting the obligation should never require transferring control of the asset. Second, a minimum arm's-length share: a defined proportion of each obligated party's incremental cover requirement should be procured from unaffiliated providers, guaranteeing a demand channel that buyer-side bargaining power cannot close; this is the same logic as existing market-making obligations and can be set at a modest initial level. Third, a central procurement backstop: where the market has not closed the shortfall by the binding notification, or where the Authority finds buyer-side concentration is suppressing independent supply, the Authority should tender for the residual cover directly and allocate its cost to obligated parties pro rata — a neutral buyer of last resort whose availability disciplines bilateral negotiations even when it never transacts.

A closing observation on complexity. The extent of safeguarding this submission has had to propose — a cover register, certification safe harbours, arm's-length procurement shares, non-discriminatory offer backstops, central procurement of last resort, universal verification — exists for one reason only: the same small set of integrated firms sits on both sides of every transaction this regime creates, as obligated buyer, dominant seller and competitor to every independent provider. In a structurally separated market, little of it would be necessary. The complexity of the safeguard architecture is itself a measure of the structural problem — a shadow price of non-separation, paid in regulatory design and enforcement risk — and it should be read as evidence when the Government weighs the two tracks described at Q4. Without these features, the reconciliation proposed at (1) is likely to show the mechanism's cost rising while the premium does not fall — because the obligation will have consolidated, rather than contested, control of the resources that set the premium.

Short-term Obligation

Q15. Do you support the introduction of a short-term obligation as a targeted backstop where a near-term winter reliability risk emerges despite the operation of the long-term obligation?

☒ Yes, I agree ☐ No, I don’t agree

Yes, as a genuinely residual backstop: trigger-based, time-limited, and switching off when conditions recover. The two layers address different failure modes (investment adequacy versus in-year fuel readiness) and the short-term layer is what makes the long-term layer credible in an unexpected bad year.

One principle should be stated in the framework to prevent the two layers from becoming duplicative: the short-term obligation is a readiness and residual mechanism, not a second procurement. In-year compliance should credit cover already verified under the long-term obligation at its full certified value, applying the same firmness standards to updated conditions, and should require additional action only on the residual between the long-term position and the requirements of the deteriorated outlook. Absent that principle — automatic recognition through a common register, consistent firmness standards across both layers, and a trigger netted for contracted-and-ready cover — obligated parties will bear duplicated compliance costs in triggered years, and those costs will pass to consumers as a second premium for a risk already paid for once.

Q16. Do you agree that the short-term obligation should apply to large generators whose contribution to annual energy supply exceeds a defined de minimis threshold? We welcome views on whether 2,500GWh of expected annual generation is the appropriate threshold.

☐ Yes, I agree ☒ I agree in part ☐ No, I don't agree

Generators are the right holders for an in-year fuel-readiness obligation, and 2,500 GWh (five-year average) is a workable starting threshold. One gap: because the obligation follows the load a generator's portfolio supports, a mid-sized generator or dedicated firming provider that has sold substantial long-duration cover contracts could be materially relevant to winter readiness while sitting below the generation threshold. A supplementary capture criterion for large sellers of cover is not needed provided verification attaches to registered instruments rather than party classes (see Q24(5)); if that principle is not adopted, a fallback criterion should capture any party whose registered cover-sale commitments exceed a defined level, regardless of its own generation volume. The threshold should be reviewed on the same cycle as the long-term threshold as new entrants scale.

Q17. Do you consider the proposed trigger – projected hydro storage crossing the Watch Curve within 90 days if it declines at an assumed rate of 18GWh/day – is appropriate? If not, what alternative trigger should be considered?

☐ Yes, I agree ☒ I agree in part ☐ No, I don't agree

The retrospective testing against 2020–2025 (capturing 2021, 2024 and 2025; not triggering in 2020, 2022 or 2023) is transparent and persuasive, and I support the backtested approach. Two refinements: (a) the 18 GWh/day decline rate is calibrated to the historical demand and generation mix; under demand growth and a changing flexible-load composition it will drift out of calibration. It should be a reviewable parameter set by the Authority within the SOSFIP-style review machinery, not hard-coded in the framework. (b) The trigger should measure residual risk, not gross hydrology: contracted and readiness-verified demand response and interruptible load should be netted off before the trigger assessment, which is exactly the direction of the contracted Electricity Risk Curve work noted at paragraph 135. Otherwise the trigger will fire in years where the demand-side cover

already in place makes the position manageable — imposing seven weeks of compliance activity to demonstrate cover the system already knew about.

Q18. Do you agree with the proposed monitoring period, notification process, three-week response window, and approach to switching off the obligation if conditions improve? If not, what alternative process would be preferable?

☒ Yes, I agree ☐ No, I don't agree

Broadly yes. One drafting point: this form describes a "three-week response window", while the discussion document (paragraph 128) specifies a three-week verification period followed by a four-week response period — seven weeks in total. The document's structure is the right one and the final design should confirm it: three weeks is enough to verify data but not to order a coal shipment or activate contracted cover at reasonable cost. The switch-off threshold (projected storage above the Watch Curve for more than 120 days at 18 GWh/day decline) is sensible, as is the principle that the obligation responds to a defined residual risk rather than compelling action once the risk has receded.

Q19. Do you agree that each obligated generator's obligation should be determined by reference to the winter load its portfolio is expected to support, including its own retail load and relevant contractual commitments? If not, what alternative basis should be used?

☐ Yes, I agree ☒ I agree in part ☐ No, I don't agree

The follow-the-exposure logic is sound: a generator should stand behind the load its portfolio is committed to serve. The design must, however, prevent the same load being claimed — or orphaned — across two parties when cover is bought and sold. Paragraph 161's observation that a generator selling cover thereby takes on additional demonstrable winter demand is exactly right and should be formalised: every registered cover transaction should automatically move the associated demonstration duty in the register, so the sum of demonstrated positions always reconciles to system load. Without that automation, portfolio accounting at the compliance date becomes contestable precisely when time is shortest.

Q20. What forms of firm energy fuel cover should qualify under the short-term obligation? In particular, how should hydro be treated, including whether, and in what circumstances, existing stored water should be capable of counting toward compliance?

Qualifying cover should be anything genuinely deliverable within the relevant winter at the standard the trigger circumstances demand: physical fuel with delivery evidence (coal or LNG ordered with confirmed shipment timing, not intentions); thermal plant demonstrated available to run; demand response in an activated or readiness-confirmed state — notice served, standby confirmed, recent activation test passed; contracted interruptible load on the same activated-readiness standard; and conservatively credited intermittent portfolios at Authority-set capacity factors. The principle in paragraph 146 is right: the regime should count accessible cover, not theoretical options — and that cuts both ways, qualifying genuinely ready demand-side resources while excluding paper arrangements of any technology.

On hydro: stored water should be capable of counting, but only above a conservative catchment-level deemed-firm line set by the Authority on hydrological assumptions consistent with the conditions that fired the trigger. The asymmetry principle at paragraph 152 is the key discipline — if thermal must demonstrate fuel, hydro must demonstrate water on equivalently sceptical assumptions, or the regime quietly converts from an ex ante obligation into an ex post sanction. I also support the safety valve at paragraph 153 (action at too high a cost need not be compelled), with one tightening: reliance on it should require a published Authority determination, not the obligated party's self-assessment, or it becomes the default exit from the obligation.

Q21. How should compliance be assessed? We welcome views on what evidence should be required, and how the Authority should approach questions of firmness, attribution, double counting, and reliance on third-party arrangements (including examples of how these matters are addressed in other jurisdictions).

Evidence should be tiered by cover type: physical (executed fuel contracts, shipment confirmations, delivery windows), contractual (the executed contract plus counterparty confirmation and a penalty indemnity), and operational (activation test results or readiness certificates for demand response and interruptible load). For third-party reliance, the document's contract-plus-indemnity approach (paragraph 163) is the right stopping point for the look-through problem — supplemented by the cover register, under which the Authority checks the counterparty's own registered compliance position rather than re-investigating underlying assets, so verification chains terminate after one hop.

Comparable machinery exists in other jurisdictions: AEMO's RERT applies readiness verification and activation testing to contracted demand-side reserves; the GB Capacity Market operates agreement management with satisfactory-performance requirements and termination for unproven capacity; PJM's Capacity Performance regime attaches non-delivery charges that make counterparty performance self-enforcing. None maps perfectly onto an energy-adequacy (rather than capacity) problem, but the evidence-and-verification toolkits translate directly.

Q22. How should the regime deal with arrangements that appear compliant in advance but fail in practice because of upstream supply failure, transport disruption, force majeure, counterparty failure, or similar events?

Same architecture as the long-term regime (Q12), compressed for the timeframe: a notified-failure re-procurement duty within a short defined cure window; penalties attaching to failure to re-procure or to misrepresentation at the compliance date, not to the supervening event; narrow force majeure that excludes ordinary commercial risk; mandatory performance warranties and penalty indemnities in material cover contracts so the delivery incentive sits with the party controlling delivery; and register status flags so deteriorating cover is visible to the Authority before winter rather than litigated after it. In-year, speed matters more than perfect attribution — the regime should prioritise getting replacement cover secured, with fault allocation resolved afterwards.

Q23. How should penalties for non-compliance be designed and calibrated so that they provide a credible incentive to secure genuine firm fuel cover without encouraging inefficient or disproportionate responses?

A published dollar-per-MWh shortfall rate, set above the cost of the marginal credible in-year action (in practice, an LNG cargo or equivalent), so that securing cover is always cheaper than paying the penalty — with escalation for repeat shortfalls and the general penalty ceiling (the higher of $10 million, three times commercial gain, or 10 per cent of turnover) available for deliberate or strategic non-compliance. The published rate again functions as a price cap on in-year cover, which matters most precisely when conditions are tight and seller leverage is greatest. The disproportionality guard should run through the Authority-determined cost-benefit route (Q20), not through penalty softening: keep the penalty predictable and let the published determination handle the rare case where compliance action would genuinely harm consumers. Penalty revenue should be recycled to consumers or backstop procurement.

Q24. What other design features, transitional arrangements, or implementation issues should be considered to ensure that the short-term obligation is workable, proportionate, cost-effective, and beneficial for consumers?

(1) One register, two layers: cover procured and verified under the long-term obligation should be recognised automatically in any short-term assessment, so the layers share a single source of truth and participants are never asked to re-prove the same cover twice in one year.

(2) Pre-qualified demand-side panels: the seven-week window is too short to originate demand-response arrangements from scratch. The Authority should maintain standing pre-qualification of interruptible-load and demand-response providers — tested, registered and priced in advance — so that in-year compliance can call on a ready panel rather than a scramble. This is also the single most effective way to develop the demand-side cover market between events.

(3) Confidentiality protocols for continuous disclosure: the verification-period disclosure obligations (paragraphs 129–130) involve commercially sensitive positions; the framework should specify what the Authority may publish, in what aggregation, and when.

(4) Post-winter public reporting: after any triggered year, the Authority should publish what actions were taken, at what cost, and with what effect on the risk curves — the in-year counterpart of the annual reconciliation proposed at Q14, and the discipline that keeps the mechanism honest over successive events.

(5) Verification should follow cover, not party class. The short-term obligation as drafted tests readiness only for cover relied on by obligated generators. Two classes of qualifying cover sit outside that net: demand response and interruptible load contracted directly by wholesale purchasers, and generation owned by parties below the 2,500 GWh threshold — including independent-retailer-owned peaking plant — whose output is registered as cover. Paragraph 141 acknowledges this gap. The resolution is to distinguish procurement duties, which should remain confined to the obligated classes on proportionality grounds, from verification, which should attach to the instrument: when the trigger fires, every arrangement registered as cover for the relevant winter must confirm readiness — plant availability and fuel evidence for generation, deemed-firmness for hydro, activation status for demand response — regardless of the size or class of the party holding it. Verification framed this way imposes no burden on small parties as such; it is a condition of an instrument participating in the compliance market, incurred only by choice and compensated by the cover revenue. It also ensures sub-threshold and demand-side cover is credited on equal firmness

standards rather than discounted for being less visible — which is a precondition for independent providers competing in the cover market at all.