The Chair

AEMC

Level 15,
60 Castlereagh Street
Sydney NSW 2000

By submission portal at www.aemc.gov.au, marked confidential

7 May 2026

Re: Submission on draft determination, Improving the NEM Access Standards Package 2 (ERC0394)

Dear Chair,

VisIRGrid Pty Ltd (the **Respondent**) makes this submission to the AEMC on the draft determination and draft rule for Improving the NEM Access Standards Package 2 (ERC0394), published 12 March 2026.

VisIRGrid supports the policy intent of Package 2. A tiered access regime for large inverter-based loads closes a real gap in Schedule 5.3, and the positions in this submission are offered to make that regime workable, not to dilute its system-security purpose.

Australia has limited operating history of how hyperscale and AI data centre loads actually behave on the power system. Those loads are not uniform: the structurally distinct UPS and backup topologies now in deployment respond very differently to disturbances. Calibrating Tier 3 obligations to that lived experience requires evidence the market is only beginning to generate. In the interim, NSPs and AEMO already hold substantial powers under the existing rules to require connection-specific performance standards.

We therefore urge the AEMC to phase implementation rather than commence the full Tier 3 reform on a compressed timetable:

- first, provide a transition and grandfathering pathway for connections already underway (Position 1) so that in-flight investment is not stranded or delayed; and
- second, direct AEMO to use a defined period of approximately 12 months to co-design the detailed Tier 3 frameworks (the topology register, the load-specific control-interaction guideline, and the cost and modelling specifications) with proponents before those obligations take effect.

This sequencing protects investment appetite at a critical moment for data centre development in the NEM while still delivering the security outcomes the rule is intended to achieve.

VisIRGrid Pty Ltd

ACN 688 675 238

Level 16,

127 Creek Street,

Brisbane QLD 4000

1.1 About VisIRGrid

VisIRGrid Pty Ltd is a co-development partner for major NEM connections, including hyperscale data centres. The Respondent is engaged in connection developments in the NEM, working with developers, OEMs and connecting NSPs on technical reference design, regulatory pathway, and connection-asset financing. The perspective offered in this submission is informed by that cross-project experience and by review of the recent international technical literature on hyperscale and AI data centre power stabilisation, including OEM architecture documentation, industry working-group output and published research.

1.2 Executive summary

The Respondent supports the AEMC's intent to clarify and strengthen the access standards for large inverter-based loads. The existing Schedule 5.3 framework is not fit for connections at the scale and inverter-content of contemporary hyperscale data centre developments, and a tiered regime closes a real gap. The international alignment with IEEE 2800-2022 and the GB Grid Code is appropriate.

The Respondent submits 13 positions in Attachment A. None is intended to dilute the system-security purpose of the rule; where a position adjusts an obligation, the adjustment is conditioned on demonstrated capability and coordination with the connecting NSP and AEMO. The substantive emphases include:

- extending the AEMC's 6-month transition window to 18 months, calibrated to the actual deployment cycle for hyperscale connection infrastructure;
- an AEMO-maintained register of generic UPS topology models reflecting the three structurally distinct architectures now in active deployment (double-conversion online UPS, flywheel-based stabiliser, STATCOM-based converter), with a pre-qualification pathway so registered models can be relied on across applications rather than re-assessed each time;
- recovery-of-active-power envelopes calibrated to redundant Tier III/IV UPS recovery profiles;
- a load-specific control-interaction framework recognising the three-layer interaction unique to AI training loads;
- coordinated treatment of load and co-located storage including islanding-capable architectures;
- coordinated, uniform Tier 1 and Tier 2 standards across NEM regions;
- a hyperscale-distinct sub-tier within Tier 3; and
- a specific load class for AI training data centres with calibrated obligations reflecting synchronised GPU step changes, virtual inertia reduction and multi-layer control.

The Respondent offers to substantiate the costs of the draft Tier 3 stack with aggregated industry benchmark data, and requests a stakeholder meeting with the AEMC project team at a mutually convenient time to walk through the positions in Attachment A and to discuss the AI training data centre load class proposed in Position 11. The Respondent thanks the AEMC for the opportunity to comment on the draft determination.

Yours faithfully,

Craig Chambers

VisIRGrid Pty Ltd

[redacted]

ATTACHMENT A: SUMMARY OF FEEDBACK

Position	Issue	Proposed amendment
1. Transition arrangements for in-flight Connection Enquiries	Draft rule Section 1.5 provides only a 6-month transitional window during which applicants close to finalising negotiations may continue under the existing framework. For Tier 3 connections, where generator-equivalent PSCAD and RMS modelling typically takes 9 to 12 months and is on the critical path, 6 months does not allow developers to absorb the new obligations without compromising in-flight project commitments. Industry deployment timelines for behind-the-meter system architectures supporting hyperscale loads (Piller SHIELDX published reference) run to 18 months from final design to power-on, well beyond the AEMC's 6-month window.	Extend the transition window from 6 months to 18 months for connections at Tier 2 or Tier 3, calibrated to a typical 9 to 12 mth generator-equivalent modelling cycle plus the time required to absorb the new obligations into in-flight design and contracting commitments. Within the transition window, applicants who lodged a Connection Enquiry under NER 5.3.4 with a TNSP before the final rule determination may elect to be assessed under the framework current at the time of lodgement. Applicants who elect to opt into the new regime can do so via a single notice to the connecting NSP and AEMO, without re-lodging.
2. Single up-front Tier 3 assessment for staged build-outs	Draft rule clause S5.3.1a(a1) classifies inverter-based loads by nameplate capacity but does not address whether a connection delivered in stages over multiple years triggers a fresh Tier 3 assessment at each commissioning increment. The default per-stage reading would impose multiple full Tier 3 assessments on a single hyperscale connection. "Large" is also not clearly defined as nameplate or contingent peak. AirTrunk's submission to the consultation paper noted that staged ride-through requirements for hyperscale loads scaling over multiple years would help enable timely development, a position the Respondent endorses.	A single up-front Tier 3 assessment, scoped to the envelope of binding development consent, covers all staged commissioning within that envelope. "Large" is defined by agreed maximum demand at the connection point (contingent peak) rather than installed nameplate, so that redundant UPS capacity installed for reliability does not itself lift a connection into a higher tier. Re-assessment is triggered only by changes to plant topology, control system configuration, or the addition of co-located plant not contemplated in the original assessment. Routine capacity increments within the assessed envelope require only a confirmation filing to the connecting NSP and AEMO, not a fresh study programme.
3. AEMO-maintained register of accepted topology models	Modelling obligations for Tier 3 loads are not specified in the draft rule itself. The AEMC has delegated modelling specification to AEMO's Power System Model Guidelines (PSMG), with PSCAD and PSSE expected at Tier 3. Three structurally distinct UPS topologies are now in active deployment for hyperscale data centres: traditional double-conversion online UPS (with NVIDIA-style rack-level power smoothing	The final determination should set the policy intent that AEMO's PSMG update process establishes a register of accepted generic UPS, rectifier and switchgear topology models, validated once with industry input, that an applicant can reference in support of a Tier 3 application. The register should explicitly accommodate the 3 structurally distinct architectures now in deployment. Vendor-specific PSCAD and RMS

	via power capping and electrolytic-capacitor energy storage[1]); flywheel-based high-inertia stabilisers with bi-directional power modules (e.g. Piller SHIELDX)[2]; and STATCOM-based static converter solutions (e.g. Siemens E-STATCOM application for AI data centres)[3]. Vendor-specific re-modelling at each connection imposes cost and lead time without proportionate security gain. IEEE 2800-2022 contemplates equivalent generic-model approaches for inverter-based generation. Generic, benchmarked industry-standard models, comparable to the composite load model in PSSE, also resolve a computational constraint: detailed vendor PSCAD and PSSE models require very short simulation time steps, and running full detailed data centre models alongside the many other user-defined models across the NEM materially slows, and can crash, wide-area simulations. Benchmarking each data centre representation against a standard model would allow connection studies to proceed on standard models, with detailed PSCAD and PSSE models reserved for the cases that genuinely require them.	modelling should be required only where the applicant departs from a registered topology, the topology is materially novel, or the connecting NSP demonstrates specific local conditions warranting deeper modelling. The register is developed jointly with industry and updated on a defined cadence. AEMO should further enable UPS suppliers and standard data centre load models to be pre-qualified ahead of any specific connection application, so that a pre-qualified model can be relied on across applications rather than re-assessed at each application, reducing developer cost and accelerating registration. Where a supplier has demonstrated a model's minimum withstand short-circuit ratio (SCR) once at pre-qualification, individual applicants relying on that registered model should not be required to re-prove the model's withstand capability at each application; the site-specific SCR actually available at the connection point remains a matter for the connection study and for the network-side obligation addressed in Position 13.
4. Recovery-of-active-power envelopes calibrated to hyperscale UPS architectures	Draft rule clause S5.3.13 sets disturbance ride-through and recovery envelopes by voltage region: Region A (90-110% nominal) requires no disconnection; Region B (70-90%) requires recovery of active power to 90-110% of pre-disturbance level within 500 milliseconds (automatic) or 1 second (minimum); Region C (0-70%) permits disconnection or active-power reduction subject to maximum current limits. The Respondent acknowledges these standards are at the same level or less onerous than the equivalent generator standards, and that loads are not required to maintain continuous uninterrupted operation (CUO). However, NVIDIA's published GB300 NVL72 documentation indicates	Maintain the 500-millisecond automatic access standard. Extend the Region B minimum access standard for active-power recovery from 1 second to 2 seconds where the connecting party demonstrates a redundant UPS architecture using staggered restart. Site-specific outcomes within the 1 to 2 second range available via negotiated access standard, subject to no person being materially worse off. IEEE 2800-2022 (clause 7.2) and the GB Grid Code (Connection Conditions CC.6.3.7), although directed to inverter-based generation rather than load, support by analogy calibrated recovery envelopes for inverter-based plant, and the Respondent submits those frameworks alongside

[1] NVIDIA, How new GB300 NVL72 features provide steady power for AI
[2] Piller, Let's talk behind-the-meter island power for AI data centres
[3] Siemens, Application of E-STATCOM for AI data centres

	rack-level power smoothing achieves up to a 30% reduction[4] in peak grid demand, leaving substantial multi-MW residual step changes that propagate upstream. Piller's published SHIELDX architecture maintains ±1% frequency under 25-30% instantaneous load steps[5] in island mode using high-inertia stabiliser plus bi-directional power modules; this island-mode result is cited for the magnitude of the step changes hyperscale plant must manage, not as a measure of grid-connected recovery performance. The Region B 1-second minimum standard is difficult to achieve for n+1 redundant Uptime Institute Tier III or Tier IV UPS where stored-energy ride-through is followed by deliberately staggered restart sequencing. Many UPS systems now in deployment are generic types that may be unable to comply with the ride-through envelope without significant balance-of-plant additions such as STATCOMs, increasing capital cost and complexity and potentially reducing reliability. Inverter-based loads elsewhere in the NEM have encountered comparable issues and required STATCOMs to comply with system standards, a direct precedent for the balance-of-plant burden a prescriptive envelope can impose on inverter-based load.	the NVIDIA and Piller technical evidence as comparators for the AEMC's working-group consideration.
5. Control-interaction evidence framework specific to non-dispatching loads	The draft rule does not contain a load-specific control-interaction clause. Schedule 5.3 clauses S5.3.3(f) and (g) constrain protection-system response to phase shifts and successive faults, and the new clause S5.3.1(a1)(2A) requires the load to provide ride-through-related information on NSP request. Together these are silent on what evidence satisfies the control-interaction obligation for a Tier 3 load that does not bid into NEMDE. AI training data centres present a three-layer control problem absent in generic IBR plant[6]: GPU software ramp profiles (millisecond), UPS or stabiliser hardware response (sub-millisecond),	The final determination should set the policy intent that AEMO publish a load-specific control-interaction guideline within 6 months of commencement, with worked examples for typical hyperscale data centre topologies including the three-layer GPU / UPS / engine-governor interaction. The minimum evidence to be specified includes: EMT studies demonstrating non-interaction with co-located IBR generation in the local zone (with particular attention to declared shortfall regions); validated PLL configuration and current-limiting parameters; and a control-block diagram lodged with AEMO. NSPs

[4] NVIDIA, How new GB300 NVL72 features provide steady power for AI
[5] Piller, Let's talk behind-the-meter island power for AI data centres
[6] Microsoft Azure, Power stabilization for AI training datacenters

	and engine governor droop (100 ms or longer lag). This three-layer interaction is qualitatively different from generic IBR control and risks inconsistent NSP application without a specific evidentiary framework, with sub-synchronous control interaction with co-located IBR generation a particular risk in zones of high IBR penetration.	should reference the AEMO guideline in connection-application assessment under clause S5.3.1(a1)(2A).
6. Coordinated treatment of load + co-located BESS, including islanding-capable architectures	Draft rule clauses S5.3.1a(a1), S5.3.13 and S5.3.14 (load) read alongside Schedule 5.2 (generator) and NER 5.20C (system strength) treat a large inverter-based load with co-located grid-forming BESS at a common connection point as separate plant. Parallel assessments duplicate modelling, complicate the connection agreement and obscure integrated control behaviour. Package 1 (final May 2025, in force) introduced the principle that access standards apply by plant type rather than registration category; this principle has not been carried into Package 2 for hybrid load-storage plant. The draft is also silent on islanding-capable architectures, which use co-located grid-forming storage (and, in some configurations, on-site rotating machinery) to disconnect from the grid during deep disturbances and continue operating from local sources, then resynchronise. Behind-the-meter island power architectures, such as the Piller SHIELDX configuration[7] combining a high-inertia flywheel stabiliser with reciprocating engines and BESS, are documented to maintain ±1% frequency under 25-30% load steps in island mode. An islanding-capable plant is functionally not connected to the grid during the deep-disturbance window; treating it under the same Region B and C obligations as a continuously connected load miscalibrates the obligations.	A single coordinated Tier 3 assessment where load and co-located storage share a common connection point, with consolidated modelling and a single connection agreement amendment. The NER 5.20C system strength obligation and Package 2 Tier 3 obligations are discharged through a single integrated submission package, with concurrent rather than sequential AEMO and NSP engagement. Where the plant is demonstrably islanding-capable through co-located grid-forming storage, anti-islanding protection coordinated with the connecting NSP, and resynchronisation controls, the Region B and C ride-through obligations under S5.3.13 apply on resynchronisation rather than on continuous active-power presence during the deep-disturbance window. The separation event itself should be subject to agreed limits on magnitude, ramp and notification, coordinated with the connecting NSP and AEMO, so that deliberate islanding of a large load is a bounded, studied contingency rather than an uncontrolled loss of load. Anti-islanding scheme, resynchronisation criteria and the protection coordination interface to be specified in the AEMO load-specific guideline proposed in Position 5.
7. Alignment with AEMO frameworks and a narrow definition of material change	Tier 3 obligations under draft clauses S5.3.13 (ride-through) and S5.3.14 (instability monitoring) sit alongside AEMO's Grid-Forming and Grid-Following Inverter Fault Current Contribution paper (3 November 2025), the 2025 Transition Plan for System Security (December 2025), the Engineering Roadmap, and the modelling specifications expected in the	The final determination should confirm that compliance with the AEMO frameworks (as amended) is sufficient evidence for the corresponding Tier 3 obligations. Where AEMO frameworks address the same matter as Schedule 5.3, AEMO framework should be deemed to satisfy the corresponding Schedule 5.3 obligation; where the AEMO frameworks

[7] Piller, Let's talk behind-the-meter island power for AI data centres

	AEMO PSMG update. Without alignment, plant could face two specifications. The draft mentions "material change" only in the context of review-extension procedure and leaves the trigger for re-modelling under a connection agreement to be set case-by-case by NSPs.	are silent, Schedule 5.3 applies in its own terms. AEMO publishes a cross-reference matrix mapping the Tier 3 clauses to the AEMO framework provisions within 3 months of commencement. "Material change" is defined narrowly as changes to plant topology, control system configuration, or addition of co-located plant not contemplated in the original assessment. Routine maintenance, vendor firmware updates and capacity-factor changes within the assessed envelope are not material change.
8. Cost evidence to support the regulatory impact assessment	The draft determination does not fully quantify the cost of Tier 3 obligations on hyperscale data centre connections. Without substantiated evidence the final rule risks under-estimating implementation cost and schedule effect on the developer community it is principally intended to regulate. The Respondent has access to aggregated industry-benchmark cost data that, subject to confidentiality clearances from contributing parties, can support the AEMC's regulatory impact assessment.	VisIRGrid offers to provide aggregated industry benchmark data in 3 categories: (a) capex uplift attributable to Tier 3 plant-design changes (per MW of installed capacity, range); (b) modelling cost (PSCAD, RMS, control-interaction studies, validation; range); and (c) schedule effect (months from Connection Enquiry to Connection Agreement under the current and proposed regimes; range). As supporting evidence, Piller's published deployment timeline indicates 18 months for behind-the-meter system architectures[8] versus more than 6 years for grid connection, a regulatory delay differential that, monetised across the hyperscale connection pipeline, supports the AEMC reflecting these inputs in the final regulatory impact assessment. The Respondent would provide these benchmarks to the AEMC.
9. Mandatory uniform Tier 1 and Tier 2 standards across the NEM	Draft rule clause S5.3.1a(a1) leaves Tier 1 and Tier 2 standards NSP-discretionary while Tier 3 is mandatory. A 30 to 100 MW inverter-based load (IBL) would face different standards in different NEM regions, undermining the predictability the rule is intended to provide. The Respondent acknowledges the Energy Networks Australia submission to the consultation paper, which argued for NSP discretion to reflect	Tier 1 and Tier 2 access standards are published directly in Schedule 5.3 by the AEMC and apply uniformly across all NSPs. NSPs retain the ability to negotiate site-specific outcomes with applicants under negotiated access standards, so flexibility is preserved where local network conditions warrant. Where an NSP and applicant cannot agree on a negotiated outcome, the disagreement is resolved under the existing

[8] Piller, Let's talk behind-the-meter island power for AI data centres

	surrounding network circumstances. Consistency at Tier 1 and Tier 2 across NEM regions is, however, of greater value to large-load developers than the regional flexibility the discretion preserves: developers compare connection options across regions and an inconsistent baseline is a barrier to that comparison. Compounding this, NSP-discretionary Tier 1 and Tier 2 standards would be applied to three structurally distinct UPS topologies now in active deployment; unless minimum standards on voltage ride-through, frequency response and fault current contribution are uniform, developers face inconsistent expectations of the same plant design across jurisdictions. National-market consistency was a foundational premise of the National Electricity Rules and the Package 1 plant-type principle.	NER dispute resolution framework, informed by AEMO technical advice. This compromise preserves NSP technical input where it adds value, delivers the developer certainty that mandatory uniformity was intended to provide, and avoids the inconsistent-application risk that pure NSP discretion creates.
10. Hyperscale-distinct sub-tier within Tier 3, defined by load characteristics	Draft rule clause S5.3.1a(a1) treats all inverter-based loads of 100 MW and above as a single Tier 3 category. AWS's submission to the consultation paper recommends establishing a separate technical category for data centre loads, on the basis that data centres present different system-security characteristics from other large IBLs. The Respondent agrees. Hyperscale data centre, hydrogen electrolyser and refinery process loads are technically distinct in 3 respects: predictable load at dispatch granularity (data centre), notwithstanding the sub-second step-change dynamics addressed in Position 11, versus variable process load (electrolyser, refinery); controllability through facility-level EMS (data centre) versus inflexible base-load (refinery); and inverter content in n+1 redundant double-conversion online UPS (data centre) versus primary process-control inverters (electrolyser, refinery). AI training workloads on GPU-based architectures (e.g. NVIDIA GB300) introduce a fourth distinguishing characteristic: synchronised power fluctuations across thousands of devices operating in lockstep, creating virtual inertia reduction qualitatively different from individually controlled inverter-based plant. Treating structurally distinct plant uniformly under a single Tier 3 risks miscalibrating obligations for all classes.	Establish a hyperscale data centre sub-tier within Tier 3, defined by load characteristics rather than industry: predictable controllable load with redundant double-conversion online UPS architecture (or equivalent flywheel-based or STATCOM-based architecture per Position 3) and EMS-coordinated dispatch. Plant satisfying the sub-tier definition includes hyperscale data centres of conventional IT cooling profile and (per Position 11) AI training data centres with GPU-synchronised step-change profiles. The sub-tier is subject to the Tier 3 obligations as calibrated by the Position 4 envelope and the Position 5 AEMO load-specific control-interaction guideline. Plant outside the sub-tier definition is subject to the general Tier 3 obligations. The sub-tier definition is technology-neutral and does not name the data centre industry; it captures any plant that satisfies the load-characteristic test.

11. Specific load class for AI training data centres	AI training data centres are demonstrably a distinct load class within Tier 3 inverter-based loads, qualitatively different from conventional data centre IT and from other large-IBL classes. Published technical evidence establishes the case: NVIDIA GB300 NVL72 documentation describes "thousands of GPUs operating in lockstep"[9] and introduces power smoothing (power capping, electrolytic-capacitor energy storage, GPU burn hardware/software) that reduces peak grid demand by up to 30%, leaving substantial multi-MW residual step changes; Piller SHIELDX behind-the-meter island power architecture documents instantaneous load steps of 25-30%[10] (compared with conventional ±5%) requiring sub-millisecond stabiliser response and ±1% frequency maintenance; ESIG Interconnection Studies Short Course (17 November 2025) and New Frontiers IBR and Data Center Modeling slides (3 December 2025)[11] collate the connection-study practice forming around AI loads internationally. AI training loads also present a three-layer control problem (GPU software ramps, UPS or stabiliser hardware response, engine governor droop) absent in generic IBR plant, and a virtual-inertia reduction effect from synchronised GPU operation that is qualitatively different from individually controlled inverter-based plant. The draft rule's uniform Tier 3 obligations were not calibrated to this profile.	AEMO is requested to publish, within 12 months of the rule commencement and with input from AI infrastructure operators (NVIDIA, Microsoft, Meta, Google, AWS) and topology vendors (Piller, Siemens, Hitachi Energy, Schneider Electric), a load-specific guideline for AI training data centres addressing: synchronised step-change limits and recovery envelopes calibrated to actual GPU step-change magnitudes; three-layer control-interaction testing covering GPU software, UPS or stabiliser hardware, and engine-governor response; the interface with islanding-capable architectures (per Position 6); and modelling representations capturing virtual inertia reduction from synchronised device operation. The Respondent submits that AI training is the most rapidly evolving load class in the NEM and warrants its own technical conversation rather than being subsumed into uniform Tier 3.
12. Integration of data centre loads into NEMDE and price formation	The draft rule does not address how large inverter-based loads, which typically do not bid into NEMDE, are represented in central dispatch for the purpose of establishing dispatch and price outcomes. As Tier 3 data centre loads scale to hundreds of MW, their treatment in NEMDE (as scheduled, semi-scheduled or non-scheduled load) affects price formation, pre-dispatch forecasting accuracy and the value of any demand flexibility they can offer. Without clarity, large controllable	The final determination should clarify how Tier 3 inverter-based loads are represented in NEMDE for dispatch and price-setting purposes, and whether large controllable data centre loads with facility-level EMS may elect to be scheduled or semi-scheduled. AEMO to confirm the treatment of these loads in central dispatch and pre-dispatch forecasting so that proponents can design EMS and market-participation arrangements with certainty. Any scheduling obligation should be calibrated to the load's actual controllability and should not

[9] NVIDIA, How new GB300 NVL72 features provide steady power for AI
[10] Piller, Let's talk behind-the-meter island power for AI data centres
[11] ESIG, Interconnection Studies Short Course – Large Load Extract (17 November 2025); ESIG, New Frontiers in IBR and Data Center Modeling Requirements, Practices, and Experience (3 December 2025)

	loads sit outside the price-formation framework even where their facility-level EMS could contribute to dispatch efficiency.	impose generator-equivalent dispatch obligations on loads that cannot practicably follow dispatch targets.
13. NSP obligation to maintain minimum short-circuit ratio at the connection point	The draft rule requires connecting loads to demonstrate ride-through and withstand capability at a minimum short-circuit ratio (SCR) but does not address the corollary question: whether, and to what extent, the connecting NSP is obliged to maintain a minimum SCR at the connection point so that a compliant data centre load can connect and remain connected. Minimum withstand SCR is a function of network strength as much as load design. Placing the full obligation on the load, without a corresponding network obligation, can require the load to fund balance of plant (such as STATCOMs) to compensate for network conditions outside its control, and interacts directly with the system-strength framework in NER 5.20C.	The final determination should clarify the allocation of responsibility for maintaining minimum SCR at the connection point between the connecting load and the NSP, consistent with the NER 5.20C system-strength framework. Where a load demonstrates withstand capability at a defined minimum SCR, the NSP should be responsible for maintaining at least that SCR at the connection point, or for transparently allocating the cost of any shortfall, so that compliant loads are not required to fund network-strength remediation that is properly a shared or network responsibility. The allocation should be addressed in the AEMO load-specific guideline proposed in Position 5.