

BLACKGRID LABS

Infrastructure Beneath the World

Helios Grid

A research proposal for coordinating electricity exchanges between Canadian grids with decisions that operators can inspect and reproduce.

BGL/HG-01 | PUBLIC RESEARCH

Published 2026-03-14 | Revised 2026-10-11

blackgridlabs.com | Part of The Veydros Group

In essence

Why we are talking about this

Electricity crosses provincial borders, but the rules and responsibilities on either side can differ. A useful coordination proposal has to explain both the physical exchange and who can approve it.

Purpose

Test whether a shared evidence record can help operators compare proposed exchanges while keeping their own dispatch authority.

What we found and focus on

- Provincial authority has to be part of the design, alongside federal responsibilities. [1]
- Ontario's renewed market illustrates that scheduling and pricing rules change over time. [2]
- A proposed exchange needs a dated record of its inputs and applicable rules.
- A line's design rating is not enough to establish what can be delivered at a given hour.
- Repeatable calculations help review a decision; they do not prove the inputs are true.
- Reliability constraints, costs and distribution of benefits need separate checks.
- The proposed layer would support operators rather than take control of dispatch.
- A bounded trial with an independent reviewer is the next useful test.

Conclusion and next steps

Start with one interface and one agreed use case. Freeze the data and rules, compare against an operator baseline, and publish the errors and limits before proposing wider use.

Full revised research

The coordination question

Helios Grid asks a practical question: can two operators review the same proposed electricity exchange and understand why it was recommended? The problem includes power flows, operating rules, contracts and responsibility for the result. Software addresses only part of it.

Section 92A establishes provincial powers over electrical generation, with additional provisions on exports and federal authority. A coordination design needs to respect that framework rather than assume a single national decision maker. [1]

What the public record establishes

Ontario's renewed electricity market launched on May 1, 2025. That is a concrete example of rules changing within an operating system. A cross-border calculation needs the rules in force for its particular interval. [2]

The Electricity Advisory Council's 2024 report addresses cooperation and investment in Canada's electricity transition. It supplies policy context, not evidence that the BGL proposal has been accepted or demonstrated. [3]

The proposed evidence record

For each candidate exchange, record the time interval, input origin, transfer limit, reservations, losses, price assumptions and approval conditions. Keep measured values separate from forecasts. If a source changes, issue a new version rather than silently replace the old value.

The recommendation should identify the constraints that bind and the assumptions that drive its benefit. A reviewer should be able to rerun the calculation using the frozen record and compare it with the operator's existing method.

Reliability and authority

PJM's operations manual is a useful reference for the detail required in scheduling and dispatch. Its rules are not automatically applicable to a Canadian interface. Local operators must specify the actual reliability and market requirements. [4]

A cost reduction in a simplified model does not establish an executable exchange. The proposal must also pass delivery, contingency and settlement checks. Operators retain responsibility for approval, and an unresolved rule should stop the affected recommendation.

How to test the proposal

Use a small, agreed historical dataset before live shadow operation. Establish a baseline and score feasibility errors, replay consistency, explanation quality and review time. Include outages, stale measurements, unavailable capacity and contradictory inputs.

An independent reviewer needs the data, model version, tolerances and complete failure record. A repeatable wrong answer is still wrong. The test should measure correctness against the agreed model and acknowledge what that model leaves out.

Conclusion and next steps

The useful research outcome is a checkable method for reviewing an exchange. This edition does not establish grid-wide savings, production readiness or superiority over existing operator tools.

Agree on one interface, collect the necessary data, define acceptance criteria and run a controlled comparison. Expand only if the evidence supports it.

Sources

Numbered citations refer to these sources. Proposed methods and next steps are BGL research proposals, not results established by the cited publications. Links checked for this revision on October 11, 2026.

[1] Department of Justice Canada: Constitution Act, 1867, section 92A

https://laws-lois.justice.gc.ca/eng/const/page-3.html

[2] IESO: Renewed Market Advisory Forum, market launched May 1, 2025

https://ieso.ca/en/Sector-Participants/Engagement-Initiatives/Engagements/Renewed-Market-Advisory-Forum

[3] Canada Electricity Advisory Council: Powering Canada, a blueprint for success (2024)

https://natural-resources.canada.ca/energy-sources/powering-canada-blueprint-success

[4] PJM: Manual 11, Energy and Ancillary Services Market Operations

https://www.pjm.com/-/media/DotCom/documents/manuals/m11.pdf



BLACKGRID LABS

Infrastructure Beneath the World

Interties as Provable Instruments

A proposal for documenting what an intertie can carry, what has already been committed, and how an exchange would be checked.

BGL/HG-02 | PUBLIC RESEARCH

Published 2026-04-22 | Revised 2026-10-11

blackgridlabs.com | Part of The Veydros Group

In essence

Why we are talking about this

A transmission link can appear underused while reliability limits, reservations or operating conditions restrict the next exchange. Capacity figures need context.

Purpose

Define an interface record that both sides can use to examine delivery, benefit and settlement.

What we found and focus on

- Hardware capability and the next available exchange are different questions.
- Limits need a time interval, operating conditions and a responsible source.
- Existing commitments must be accounted for before new capacity is allocated.
- An interface should be studied on its own terms, not as a generic national seam.
- Metering and schedule records need a common time basis.
- Benefit estimates depend on the chosen baseline and should expose it.
- Settlement rules and decision authority matter alongside the physical model. [1]
- No public dataset in this edition establishes utility fault or recoverable savings.

Conclusion and next steps

Build one interface record with the operators. Reconcile schedules, measurements and limits, then test whether another analyst can reproduce the same exchange assessment.

Full revised research

An intertie as an operating agreement

An intertie connects systems whose operators may use different scheduling, pricing and approval processes. Reviewing an exchange means reviewing the equipment and the agreement under which it runs. A single capacity number cannot describe both.

This paper proposes an evidence record for that review. It does not report a completed national inventory or an operator-approved ledger.

Capacity with a timestamp

Record design capability separately from the applicable operating limit. The latter needs the interval, direction, relevant outages, contingency assumptions and reservations. When a limit changes, preserve the previous value and the reason.

An observed flow below a published limit is a starting point for investigation. It is not enough to call the difference wasted capacity. The record may lack a security constraint, a commercial restriction or an unavailable supply source.

Four checks on an exchange

First, reconcile metering and schedules on a common clock. Second, identify the physical and contractual constraints. Third, compare the proposed exchange with an explicit baseline. Fourth, calculate how the accepted exchange would settle.

PJM Manual 11 illustrates the operational specificity needed in market scheduling. It is a reference for method, not a substitute for the rules of the interface under study. [1]

Benefits that a reviewer can challenge

A claimed benefit needs the same demand, equipment availability and external conditions in both cases. Otherwise the comparison may credit the exchange for an unrelated change. Separate total system benefit from the allocation of that benefit between parties.

Publish sensitivity to losses, prices and unavailable capacity. A low-cost scenario can become infeasible under an outage. That should appear as a failed case, not disappear into an average.

Governance and missing information

The constitutional framework and the Electricity Advisory Council's policy work supply context for coordination. Neither validates a particular commercial arrangement or a BGL calculation. Actual authority and agreements must be established with the relevant parties. [2][3]

Some required information will be commercially sensitive. The record needs a controlled review process, with access appropriate to each party. Confidentiality should not be presented as evidence that an unpublished calculation is correct.

Conclusion and next steps

The immediate goal is modest: make one exchange assessment easier to inspect and dispute. A shared record may help, but the operators have to test whether it captures the details that drive their decisions.

Select one interface, agree on the record format and reconcile a sample period. Document every gap. Only then run benefit comparisons and consider a broader inventory.

Sources

Numbered citations refer to these sources. Proposed methods and next steps are BGL research proposals, not results established by the cited publications. Links checked for this revision on October 11, 2026.

[1] PJM: Manual 11, Energy and Ancillary Services Market Operations

https://www.pjm.com/-/media/DotCom/documents/manuals/m11.pdf

[2] Department of Justice Canada: Constitution Act, 1867, section 92A

https://laws-lois.justice.gc.ca/eng/const/page-3.html

[3] Canada Electricity Advisory Council: Powering Canada, a blueprint for success (2024)

https://natural-resources.canada.ca/energy-sources/powering-canada-blueprint-success



BLACKGRID LABS

Infrastructure Beneath the World

Deterministic Dispatch at the Seams: Formulation and Evidence Architecture

A proposed formulation and verification plan for repeatable cross-border exchange recommendations.

BGL/HG-03 | PUBLIC RESEARCH

Published 2026-05-30 | Revised 2026-10-11

blackgridlabs.com | Part of The Veydros Group

In essence

Why we are talking about this

A solver can return a repeatable answer while using a bad constraint or incomplete data. Dispatch research needs checks on the model, the inputs and the result.

Purpose

Specify a coordination model and an independent verification process that can be tested against a frozen benchmark.

What we found and focus on

- Optimization is already part of electricity-market operations. [1]
- The objective and constraints must be agreed before a result can be judged.
- Freeze input versions, units, time intervals and source identities.
- Control solver versions, numerical tolerances and tie-breaking when testing replay.
- A separate checker should test constraints and the claimed objective value.
- Formal verification can also apply to some neural-network power-system models. [2]
- A proof is limited to its assumptions and cannot certify unmeasured physical conditions.
- This edition makes no claim of a completed BGL benchmark or production dispatch.

Conclusion and next steps

Publish a small formulation, dataset and checker. Test normal cases and planted defects, then compare correctness and runtime with an established optimization baseline.

Full revised research

Scope of the formulation

The proposed system recommends exchanges between operators. It does not replace their dispatch engines. The research task is to define an exchange model precise enough for a second implementation to inspect.

PJM Manual 11 documents optimization within electricity-market operations. BGL's contribution would therefore have to be demonstrated in the coordination and evidence process, rather than claimed from the use of optimization alone. [1]

A minimal model

Let x represent exchange quantities over defined time intervals. A first model minimizes an agreed cost c^Tx subject to balance equations Ax = b, operating inequalities Gx <= h, and bounds l <= x <= u. This is a proposed starting point, not a complete grid model.

Each matrix row needs a meaning, unit, source and validity interval. Losses, reserve needs and contingency conditions require explicit treatment. Add integer variables only where discrete operating decisions require them, and state which nonlinear effects the model approximates.

Repeatability by construction

A reproducible run needs more than a fixed random seed. Preserve the dataset, model, solver build, numerical tolerances, ordering and tie-breaking rules. Define whether the acceptance test requires byte-identical output or equivalent decisions within a stated tolerance.

Record infeasibility and timeouts as outcomes. An implementation that substitutes a convenient answer when a constraint cannot be resolved undermines the entire review process.

Independent checking

The checker should recompute constraint residuals and the reported objective from the frozen inputs. For a linear program, an appropriate certificate can support an optimality claim. For a mixed-integer problem, report the bound and optimality gap rather than imply certainty from a feasible solution.

Verification is not exclusive to deterministic software. Venzke and Chatzivasileiadis demonstrate formal methods for specified neural-network power-system applications. Their results support a bounded verification approach, not a universal guarantee for learned systems or the BGL proposal. [2]

What the benchmark must include

Publish cases with incompatible units, stale limits, missing reservations, conflicting rules and infeasible exchanges. Compare the proposed checker with an independently implemented baseline. Separate parsing failures, model failures and solver failures.

Measure feasibility, objective agreement, repeatability and runtime on stated hardware. The earlier internal run names and latency claims do not supply a reproducible benchmark and are not evidence in this revised edition.

Conclusion and next steps

The end goal is a coordination result whose assumptions and calculations another party can inspect. Repeatability supports that goal but does not establish real-world correctness by itself.

Release a bounded test specification, obtain an independent implementation and publish disagreements. Operator review of the physical assumptions is necessary before any live recommendation is considered.

Sources

Numbered citations refer to these sources. Proposed methods and next steps are BGL research proposals, not results established by the cited publications. Links checked for this revision on October 11, 2026.

[1] PJM: Manual 11, Energy and Ancillary Services Market Operations

https://www.pjm.com/-/media/DotCom/documents/manuals/m11.pdf

[2] Venzke and Chatzivasileiadis: Verification of Neural Network Behaviour, Formal Guarantees for Power System Applications (2020)

https://arxiv.org/abs/1910.01624v4



BLACKGRID LABS

Infrastructure Beneath the World

The Demand Shock: Data Centres, Electrification, and the Coordination Gap

Research on separating announced electricity demand from evidenced commitments and testing capacity under different load scenarios.

BGL/HG-04 | PUBLIC RESEARCH

Published 2026-07-08 | Revised 2026-10-11

blackgridlabs.com | Part of The Veydros Group

In essence

Why we are talking about this

New compute demand and electrification make planning more uncertain. Adding project announcements together can hide duplication and differences in how likely, flexible or time-sensitive each load is.

Purpose

Develop a demand record that distinguishes evidence from forecasts and supports transparent capacity decisions.

What we found and focus on

- The IEA projects substantial growth in data-centre electricity demand; its outlook is a scenario, not a list of committed Canadian loads. [1]
- An announcement, queue application and signed connection agreement are different evidence classes.
- The same project may appear in more than one record and needs an identity check.
- Annual energy and demand at the constrained hour need separate accounting.
- Electric vehicles, heating, industry and data centres need different load assumptions.
- Flexibility should be counted only under stated operating and contractual conditions.
- Forecast uncertainty remains even when the accounting is repeatable.
- This edition provides no verified national total or estimate of phantom load.

Conclusion and next steps

Design a small demand register, reconcile it with its sources, and test capacity decisions across explicit weather, timing and project-completion scenarios.

Full revised research

What is changing

The IEA's 2025 Energy and AI report projects growth in electricity use by data centres. It explains why planners need to examine compute demand, but its global scenarios do not establish the load at a Canadian substation. [1]

This paper focuses on the accounting needed before local allocation decisions. It does not claim that a private BGL register has measured the Canadian queue or established a national capacity shortfall.

A record of commitments

For each project, record the applicant, location, proposed size, timing, source date and evidence status. Keep announcements, applications, agreements, construction and operating loads in separate categories. A project moving between stages should keep its history.

Check for duplicate names, corporate affiliates and alternative sites. A duplicated entry should be resolved with evidence rather than a guess. If uncertainty remains, retain it and show how it changes the aggregate.

Energy, peak and flexibility

A project's annual consumption is not its contribution to the constrained hour. Planning needs a load shape and the conditions under which that shape changes. Seasonal heating demand, industrial shifts and compute workloads require different assumptions.

Flexible demand also needs a definition: how much can be reduced, for how long, on what notice, and at what cost? A proposed flexible workload should not be treated as dependable relief without an enforceable arrangement and a performance test.

Forecasts remain necessary

Evidence about present commitments does not determine future demand. Completion dates, utilization and weather remain uncertain. A governed register can make those assumptions visible while forecasts and scenarios estimate what could happen.

The Electricity Advisory Council's 2024 report supplies a broader planning context. It does not prove that ledger software eliminates forecasting uncertainty or settles competing claims on capacity. [2]

A testable research method

Begin with a limited region and publicly traceable projects. Have another analyst classify the same records independently. Measure disagreements and the frequency of duplicates before using the register for planning.

Run several demand scenarios with identical network assumptions. Report which decisions change when a project is delayed, removed, duplicated or less flexible than expected. Keep insufficient evidence visible instead of burying it in a single forecast.

Conclusion and next steps

Better evidence about commitments can improve a capacity discussion. It will not remove the need for forecasts, engineering studies or a legitimate allocation policy.

Publish the classification rules and a small example register. Reconcile it independently, then test whether it changes a real planning decision in a useful and defensible way.

Sources

Numbered citations refer to these sources. Proposed methods and next steps are BGL research proposals, not results established by the cited publications. Links checked for this revision on October 11, 2026.

[1] IEA: Energy and AI, executive summary (2025)

https://www.iea.org/reports/energy-and-ai/executive-summary

[2] Canada Electricity Advisory Council: Powering Canada, a blueprint for success (2024)

https://natural-resources.canada.ca/energy-sources/powering-canada-blueprint-success

BLACKGRID LABS

Infrastructure Beneath the World

Atlantic Corridor: Initial Research

An initial feasibility question: can Atlantic Canadian power, cooling and connectivity support a competitive compute service?

BGL/AC-01 | PUBLIC RESEARCH

Published 2026-06-19 | Revised 2026-10-11

blackgridlabs.com | Part of The Veydros Group

In essence

Why we are talking about this

Atlantic Canada has energy and connectivity assets worth studying. Geography alone does not establish an economical, reliable data-centre location.

Purpose

Determine which combinations of deliverable power, network service and operating cost merit a detailed feasibility study.

What we found and focus on

- Provincial electricity systems need separate assessments; Nova Scotia's CER profile is one starting point. [1]
- Power must be available under workable connection and supply terms.
- Cold weather may reduce cooling demand, but savings require a facility-specific model.
- Storm exposure, backup power and maintenance can change the economics.
- Fibre routes need measured latency, diversity and contracted capacity.
- The product would be compute service, with its own customer and utilization risks.
- Canada's sovereign-compute policy supplies context, not a customer commitment. [2]
- Private sites, counterparties and follow-on papers remain restricted.

Conclusion and next steps

Compare representative scenarios using current supplier evidence. Obtain power terms, measured network performance and a complete cost model before naming a preferred site or claiming an advantage.

Full revised research

The feasibility question

The Atlantic Corridor concept combines power supply, cooling and network access to deliver compute services from Atlantic Canada. Its value depends on the delivered service, not simply on the region's proximity to Europe or its climate.

This is an initial research proposal. It does not establish that Atlantic Canada has the continent's cheapest cooling, a proven latency advantage or bankable power at a particular site.

Power first

The CER's Nova Scotia profile documents the province's electricity system and energy context. It is useful background, but a provincial summary cannot determine the capacity available to a new facility. [1]

A candidate needs a connection study, a delivery schedule and supply terms that cover both normal operation and interruptions. Assess renewable output, firming, network charges and backup fuel together. Announced generation is not contracted firm capacity.

Cooling and resilience

Cold weather can offer an opportunity for cooling, but the size of that opportunity depends on equipment, temperature and humidity limits, load density and the cooling design. Build an hourly model and compare it with a realistic alternative location.

Include flooding, wind, salt exposure, service access and backup requirements. A lower cooling bill can be outweighed by resilience or maintenance costs. The proposed research should measure the combined effect rather than assign a regional superlative.

Network and customer requirements

A fibre landing nearby is not enough. Request service quotes, measure round-trip latency along the actual route, and verify physical diversity. A second provider does not ensure an independent path if both services share a conduit.

Define the workload before ranking sites. Batch compute, interactive services and distributed training have different network requirements. Compare revenue under credible utilization levels and include the risk of an unfilled facility.

Policy and project gates

Canada's December 2024 sovereign-compute announcement is relevant policy context. It does not establish an Atlantic buyer, a procurement award or financing for this proposal. [2] The IEA's energy analysis gives a broader demand context, with projections that remain uncertain. [3]

A detailed feasibility study should require credible power terms, measured connectivity, a resilience plan and customer evidence. Stop or redesign if the combined cost cannot compete under conservative assumptions. Keep commercial and partner information within its existing access controls.

Conclusion and next steps

The concept is worth testing where the power and network evidence intersect. The current public edition does not select a site or claim completed commercial validation.

Build a transparent comparison of representative cases. Seek independently reviewable measurements and quotes. The five restricted follow-on papers and the separate portal document remain private.

Sources

Numbered citations refer to these sources. Proposed methods and next steps are BGL research proposals, not results established by the cited publications. Links checked for this revision on October 11, 2026.

[1] Canada Energy Regulator: Nova Scotia energy profile

https://www.cer-rec.gc.ca/en/data-analysis/energy-markets/province-territory-energy-profiles/nova-scotia.html

[2] ISED: Canadian Sovereign AI Compute Strategy, announcement of December 5, 2024

https://www.canada.ca/en/innovation-science-economic-development/news/2024/12/canada-to-drive-billions-in-investments-to-build-domestic-ai-compute-capacity-at-home.html

[3] IEA: Energy and AI, executive summary (2025)

https://www.iea.org/reports/energy-and-ai/executive-summary



BLACKGRID LABS

Infrastructure Beneath the World

On the Limits of Statistical Language Systems

A practical boundary for language models: useful assistance, with independently checked evidence and rules for consequential conclusions.

BGL/DI-R1 | PUBLIC RESEARCH

Published 2026-08-02 | Revised 2026-10-11

blackgridlabs.com | Part of The Veydros Group

In essence

Why we are talking about this

A fluent answer can still be wrong. Institutions need to distinguish useful generated text from a conclusion they can defend with evidence.

Purpose

Define where language-model assistance fits and what checks a consequential decision needs.

What we found and focus on

- The Transformer is an architecture; it does not itself guarantee factual accuracy. [1]
- GPT-4's technical report describes next-token pretraining and subsequent alignment. [2]
- TruthfulQA documents false answers on a defined benchmark of earlier models. [3]
- Those results do not prove that every present or future model has the same limitations.
- A generated explanation is not automatically a valid derivation.
- Retrieval, tools and human review can help, but need evaluation on the actual task.
- A deterministic verifier can also be wrong if its rules or inputs are wrong.
- BGL's proposed boundary puts accountable checks between model output and consequential decisions.

Conclusion and next steps

Specify the task, evidence requirements and verifier. Evaluate the complete workflow on held-out cases, including misleading sources and correct-looking but unsupported answers.

Full revised research

The question this brief addresses

Language models can help draft, extract and organize information. The question here is what must happen before an institution treats an output as a consequential conclusion. Readability and confidence are not sufficient evidence.

This edition states a proposed workflow boundary. It does not disclose the private initiative's methods, partners or program, and it does not claim results from a sealed BGL evaluation.

What the cited literature shows

Attention Is All You Need describes the Transformer architecture. The GPT-4 technical report describes next-token pretraining followed by alignment, while also reporting capabilities and limitations. Neither supplies a general proof that a generated answer is correct. [1][2]

TruthfulQA measured responses to questions designed to expose common false beliefs in a specified set of earlier models. That finding is evidence about a benchmark and those models, not a universal impossibility theorem. New systems and workflows need their own evaluation. [3]

Explanations and evidence

An explanation can sound like a derivation without establishing that each step follows from an admitted premise. For a consequential conclusion, identify the source facts, applicable rules and checks that connect them to the result.

A citation should support the associated claim. A real link attached to an unsupported sentence is still an unsupported sentence. Check whether the source says what the answer claims, whether it is current enough and whether contradictory evidence has been omitted.

A proposed operating boundary

Use model output as a candidate where that helps the task. Validate extracted facts against their sources and test proposed calculations or code with an appropriate checker. Keep a named owner responsible for the final decision.

Independent checking can include deterministic rules, formal verification or human review, depending on the task. Research on power-system neural networks shows that formal guarantees can be developed within specified models and domains. That supports carefully bounded verification rather than a blanket exclusion of learned tools. [4]

How to evaluate the whole workflow

Freeze the task, dataset and scoring rules before testing. Include unsupported answers, contradictory sources, stale documents and cases that require refusal. Score the final accepted conclusions, not just the model's raw responses.

Test the verifier too. It may accept a valid calculation built from a false premise or reject a correct answer for the wrong reason. Record human overrides and review whether they improve outcomes. Repeatability is useful evidence about consistency, not a substitute for correctness.

Conclusion and next steps

BGL's proposed position is that consequential conclusions need accountable evidence and validation. Language models can contribute where the full workflow meets that standard.

Publish a bounded evaluation protocol and independent results before making broader performance claims. Keep the initiative's restricted research private.

Sources

Numbered citations refer to these sources. Proposed methods and next steps are BGL research proposals, not results established by the cited publications. Links checked for this revision on October 11, 2026.

[1] Vaswani et al.: Attention Is All You Need (2017)

https://arxiv.org/abs/1706.03762

[2] OpenAI: GPT-4 Technical Report (2024 revision)

https://arxiv.org/abs/2303.08774v6

[3] Lin, Hilton and Evans: TruthfulQA, Measuring How Models Mimic Human Falsehoods (2022)

https://arxiv.org/abs/2109.07958v2

[4] Venzke and Chatzivasileiadis: Verification of Neural Network Behaviour, Formal Guarantees for Power System Applications (2020)

https://arxiv.org/abs/1910.01624v4