

Interconnection studies that took months now run in minutes

A major North American utility ran interconnection impact studies on ThinkLabs' physics-informed AI digital twin, evaluating 8760 hours, contingency cases, and load flexibility options for each request inside its own secure Azure environment.

MAJOR NORTH AMERICAN INVESTOR-OWNED UTILITY / TRANSMISSION AND DISTRIBUTION NETWORKS

~4,000

BUSES IN THE
PLANNING MODEL

~20

LARGE LOAD
PROJECTS EVALUATED

~4 GWs

OF LARGE LOAD FLEXIBILITY
AND INTERCONNECTION STUDIES

01 / THE CHALLENGE

More requests than existing tools and processes could study

Demand arrived faster than the utility could study it. Clearing one large load request means testing 8760 hours against a dozen network states, ~100,000 AC power flow solves per year for a single candidate, and commercial planning tools are built for snapshots rather than years. Transmission planners ran a peak-hour case and set curtailment by trial and error, which stretched each study into months.

A peak-hour study fails in both directions. It rejects capacity that is deliverable for most of the year, and it can miss the hour the network is most constrained.

Cost of delay

Long backlogs and study cycle times mean developers without an energization date go elsewhere.

Cost of conservatism

Capacity stranded by worst-case assumptions is revenue the utility does not earn and cost the ratepayer carries.

Cost the hiring gap

The engineering hours required, and the retirement curve behind them, mean the problem cannot be hired away.

02 / THE ENGAGEMENT

What was deployed, and how

CLIENT PROFILE

Major North American investor-owned utility

- Millions of customers
- Separate T&D planning teams
- Significant data center growth

Engagement ran against the utility's own planning cases, covering transmission flexibility with N-1 contingency screening.

DURATION

14-week engagement focused on transmission

ENVIRONMENT

The utility's own Azure GPU environment, Kubernetes orchestrated

INTEGRATIONS

Trained on historical and simulated grid datasets

ECOSYSTEM

Partnered with NVIDIA, Microsoft and WWT



Week 4

On-premises data and infrastructure preparation



Week 8

AI training on utility specific datasets



Week 10

Large load interconnection and flexibility studies



Week 12

Results validation



Week 14

Next phase planning

1

8760 power flow

Full-year timeseries rather than a peak snapshot, because the most limiting condition often falls outside the peak-load hour.

2

Contingency screening at volume

Every hour evaluated against 100,000s of scenarios, with thermal and voltage violations classified automatically rather than post-processed.

3

Load flexibility generation

The twin returns an hourly, secure curtailment schedule, so a request that fails at peak can still connect on flexible terms.

4

Capacity mining

Hour-by-hour analysis across the region surfaced headroom that worst-case screening had hidden, on lines the utility already owned.

Every result defensible in technical review because each traces back to a solved network state.

3 min

8760 ANALYSIS
RUN ON 10 CONTINGENCIES

>99%

VIOLATION RECALL,
THERMAL AND VOLTAGE

100s of MWs

CAPACITY
IDENTIFIED

DIMENSION	BEFORE	AFTER
Time to run a study	5 days on 4 CPU workers	3 min on 4 GPUs
Analysis horizon	Peak-hour snapshot	100,000s of scenarios
Contingency review	Manual post-processing	Classified automatically, >99% recall
Curtailment level	Set by trial and error	Hourly secure schedule, generated
Solution set	Wires upgrades	Wires and flexibility, analyzed together

Planners get a model they trust, and AI agents that run the studies end to end at the pace the queue sets.

The digital twin deploys inside your security boundary, on infrastructure you already trust.

- Runs where your data lives**
On-premise or in your own cloud tenant, GPU accelerated. Network models never leave your security boundary.
- Planners work in a UI with a Chat Interface**
Study setup, results and reports without touching a command line.
- Systems call the API**
REST and CLI endpoints, orchestrated on Kubernetes, so results land in the tools your teams already run.

What this looks like on your network. Your model in, digital twin trained in your environment, accuracy benchmarked against your classical solver, hourly hosting capacity out. Weeks from environment setup to first study.

PARTNERS

Scope a deployment

thinklabs.ai / contact@thinklabs.ai