

Data Centers at the Door: Large-Load Interconnection

Honest queue totals, feasibility at every ramp step, and connection options compared on one scale

PRODUCT

Grid Data Enhanced Analytics

AUDIENCE

Planning and interconnection engineers, utility executives

A new kind of queue

For years, load growth at most utilities was slow enough to plan with a spreadsheet and a long memory. Data centers changed that. A single campus can ask for as much load as a small city, on a ramp schedule that stretches across several years, at a site chosen for land, fiber, and water rather than for grid capacity.

The requests arrive faster than planning teams can study them, and every part of the utility wants a different answer. Executives want to know how much load is really coming. Planners want to know whether a given request can be served, and when. Interconnection engineers want to know what it would take, and what it would cost. All three questions depend on the same network, and the answers have to agree.

Why it is hard

Queue totals are inflated. Developers often file for the same project at more than one site, or with more than one utility, to find the fastest connection. Some requests are placeholders that will never be built. Add them all up and the gross total overstates what will arrive. Load forecasts, resource plans, and rate discussions built on that total inherit the error. Quietly dropping suspected duplicates is no better, because nobody can prove which ones will fall away.

Final load is the wrong question on its own. A data center rarely switches on at full size. It ramps in phases as buildings and computing halls come online. A site that cannot take the final load may be able to take the early phases today, with upgrades timed for the later ones. Studying only the final number says "no" too often, or says "yes" years too late.

There are many ways to connect. A request can be served by reconductoring an existing feeder, adding a new feeder, adding substation transformer capacity, moving the point of connection, serving at a higher voltage, or offering a flexible service. Each option admits a different amount of load at a different cost. Comparing them by hand, for every request, does not keep pace with the queue.

The point of connection has its own risks. Large computing loads bring specific stress to the local system. Engineers have to know whether the available fault current at the point of connection stays within the ratings of the equipment there. They have to know how large a voltage change neighbors will see when the facility's load

steps on or off, including when its uninterruptible power supplies transfer during a disturbance. Reliability organizations have started to study cases where large computing loads disconnect together during a grid event. And because data centers usually want redundant supply, engineers have to know whether an alternate path can carry the load if the normal supply is lost, the classic N-1 question.

Principles of a good approach

Treat the queue as a register. Each request carries its site, its ramp schedule, its status, and its relationship to other requests. The queue becomes something the utility can reason about, rather than a list it can only count.

Report gross and net, side by side. Flag likely duplicate requests and show both totals: everything filed, and the total after netting out likely duplicates. Label which is which. Let planners confirm or clear each flag, and never drop a request silently.

Study every ramp step. Check feasibility at each step of the ramp, not only at the final load, and name the binding constraint at each step so the team knows exactly what to fix and when.

Compare options on one scale. Build a ladder of connection options and rank them by cost per megawatt admitted. That puts a cheap option that admits little load and an expensive option that admits a lot on the same footing.

Screen the point of connection early. Check fault duty, voltage step, and N-1 alternate supply before a request advances, not after a service agreement is drafted.

Refuse with reasons. When a scenario cannot hold, the tool should say so and say why, rather than returning a number that looks like an answer.

What it looks like in practice

Grid Data Enhanced Analytics (GDEA) includes a complete large-load interconnection workflow built for the data-center era. The queue lives in a register with duplicate-candidate netting and honest gross and net totals. Feasibility is checked for the point load at each ramp step. Connection options sit on a ladder ranked by cost per megawatt admitted. Fault duty, UPS voltage step, and N-1 alternate-supply checks run at the point of connection.

All of it runs on the same joined picture the rest of GDEA uses: the utility's GIS connectivity model, live SCADA telemetry, and AMI data, with real power-flow solvers behind it. The scenario sandbox solves the base case and the proposed change together, so the difference a planner sees is the change and nothing else.

ILLUSTRATIVE EXAMPLE (HYPOTHETICAL)

The following scenario is hypothetical. A mid-sized utility has a handful of large-load requests near one substation. Two come from related developers, on nearby parcels, with similar ramp schedules. The register flags them as likely duplicates. The executive summary shows the gross total with all requests included and the net total with the pair counted once, each clearly labeled. After a call with the developer, the planner confirms one of the pair will be withdrawn.

For the request that remains, the feasibility view shows the first ramp step fits on existing feeders today. The second step is bound by the substation transformer. The final step needs a new supply. The connection ladder ranks the options, and a transformer upgrade timed for the second step comes out well ahead of a new

substation on cost per megawatt admitted. The point-of-connection screens show fault duty stays within ratings after the upgrade, the voltage step is acceptable for the early phases, and the alternate feeder can carry the first step but not the second. That last finding becomes a condition in the discussion with the developer, before anything is signed.

The executive, the planner, and the interconnection engineer are all looking at the same numbers.

Questions to ask any vendor

- Does the tool show gross and net queue totals, clearly labeled, and how are likely duplicates flagged and confirmed?
- Can it check feasibility at each ramp step and name the binding constraint at each one?
- Does it compare connection options on a common measure, such as cost per megawatt admitted?
- Does it screen fault duty, voltage step, and N-1 alternate supply at the point of connection?
- Is it studying your as-built network model with real power flow, or a simplified capacity table?
- When a scenario cannot hold, does it explain why?
- Can executives, planners, and interconnection engineers work from the same results?

Closing

The data-center queue will not wait for a planning cycle. Utilities that can separate real load from inflated totals, study each ramp step, and compare options on one scale will make better commitments, sooner, and defend them with evidence.

Learn more about Grid Data Enhanced Analytics at perinimble.com/grid-data-enhanced-analytics/, or talk with our team at perinimble.com/contact/.

ABOUT PERINIMBLE

PeriNimble builds mission-critical grid data software for energy and utility companies. Founded in 2022 and headquartered in Bloomington, Minnesota, we support more than 15 million meters and over 1 billion interval reads per day in production.

Built for systems that can't afford failure. · perinimble.com · perinimble.com/contact