AEIC Large-Load Forecasting Survey —

What the Cohort Told Us

Summary of responses as of June 25, 2026. Thirteen responses from eleven utilities. Individual answers are held in confidence; everything below is aggregated, with illustrative quotes anonymized.

Participating utilities: APS, Alabama Power, ComEd, Con Edison, Dominion Energy, DTE Energy, Duquesne Light, Eversource, Oncor, Tampa Electric, and TVA.

Headline findings

- **Large load is now a top growth driver, not an edge case.** Six of eleven utilities said large loads are at least a quarter of projected growth over their planning horizon, and four said more than half.
- **There is no shared definition of "large."** Stated thresholds run from 1 MW to 100 MW, clustering loosely at 5, 25, and 50 MW. Several utilities have no firm cutoff, and a few note their ISO or RTO is coalescing around roughly 50 MW.
- **Almost everyone keeps large loads in a side tracker.** The dominant method is a separate spreadsheet or tracker, modeled project by project, not inside the main end-use or econometric model.
- **The hard problem is the same everywhere: will it actually show up, and when.** Materialization, energization timing, and ramp uncertainty came up again and again as the biggest challenge.
- **Most utilities refuse to forecast speculative load.** Eight of thirteen do not include loads that have not formally requested service. The gate is almost always a signed contract or commitment.
- **The cohort is converging on scenarios but is not there yet.** Roughly half still produce a single deterministic forecast; the rest layer in scenarios, ranges, or confidence intervals, and several are mid-transition.

What counts as a "large load"

No common threshold. The range:

- **1 MW** at the low end (one utility flags every addition of 1 MW or more, then assigns a certainty level).
- **5 MW** for distribution planning at several utilities, tied to available transformer sizes.

- **10 to 30 MW** is the most common working band (one utility uses 10 MW at 12 kV and 20 MW at 25 kV; others cite 25 or 30 MW).
- **50 to 100 MW** for transmission and RTO visibility (one reports >50 MW to its RTO; one runs a 100 MW large-load queue while its tariff defines large load above 25 MW).

A few have no firm definition yet and are leaning toward the ~50 MW threshold their RTO/ISO is settling on.

The pipeline

- **Data centers dominate** every pipeline. Behind them, in rough order of frequency: EV / fleet / transit electrification, building-electrification aggregations, manufacturing and semiconductor fabs, crypto and bitcoin mining, then a long tail (hydrogen and electrolyzers, oil and gas electrification, ports, controlled-environment agriculture, stadiums and transit).
- **Tenure varies widely.** Two utilities have only two to three years of large-load experience; one has fifteen. Data centers specifically surged for most respondents around 2023 to 2024.
- **Scale already connected is enormous and uneven.** Reported aggregate interconnected capacity ranges from a couple hundred MVA to multiple gigawatts, with the largest respondents citing roughly 6 to 10 GW already on their systems.
- **Share of future growth (eleven answered):** more than 50 percent (4), 25 to 50 percent (2), 10 to 25 percent (3), under 10 percent (2).

How utilities build the forecast

- **Who owns it:** most often a joint effort across load forecasting, transmission or distribution planning, and economic development or interconnection. A few seat it squarely in load forecasting; one has a dedicated data-center practice.
- **Method and tools:** a separate large-load tracker or spreadsheet is by far the most common, kept independent of the main model. A handful use a dedicated tool (network modeling software, an internal web-based projection tool, a dashboard). One layers quantile regression on historical demand and ramp data.
- **Horizon:** from 5 years to 30 years, commonly 10 to 20. Several run a long-term view annually and a near-term view monthly.
- **Spatial granularity:** substation is the most common anchor, followed by individual project or customer, then system, feeder, and transmission delivery points. Most use a mix depending on data availability.
- **Refresh cadence:** annual is the norm, with several on a twice-yearly cycle. Off-cycle updates are triggered by a newly signed contract, an overloaded substation, or a material change to a very large customer.

- **Deterministic vs. scenario (thirteen answered):** deterministic only (4), scenario or probabilistic ranges (4), deterministic plus scenarios (3), in development (1), one blank. Confidence intervals and low / expected / high cases are the common scenario forms.
- **Representation:** large loads are almost always carried **individually by project**. Some also fold them in as an aggregate adder; one carries a single large industrial load exogenously.

The speculative-load problem

This is where practice diverges most.

- **Include loads not yet formally requesting service? Mostly no.** Eight of thirteen do not. Four take a conditional approach, and one includes them through an economic-development pipeline.
- **Signals used to judge how real a prospective load is:** developer conversations lead, followed by land acquisition, zoning and permitting activity, applications, financing or deposits, construction progress, and news. No one relies on a single signal; it is a combination that raises confidence.
- **How a project gets gated in:** the common gate is a signed contract, load commitment agreement, or service agreement. Variations include a completed planning study plus financial backing, a confidence threshold (one utility uses greater than 75 percent likelihood), and a stage-gate maturity process. One standout uses a **weighted probability matrix**: economic-development win probability, queue milestones, contract-demand realization rates, and project-actualization rates, each weighted and multiplied against requested size.
- **Lead times** on service requests run from about 12 months to 60-plus months, clustering at 18 to 36 months.

Ramp, profiles, and behind-the-meter

- **Ramp-up uncertainty** is handled case by case. The common approach: start with the customer-provided ramp, then adjust it toward what similar customers have actually done, and apply a probability of delay. One utility requires a fixed multi-year ramp commitment and will not study alternatives; another does not yet discount customer ramp at all. Several explicitly named ramp timelines as something they want to learn from peers.
- **Co-located generation and behind-the-meter** treatment is mixed. Some net it out of the load forecast, some handle it in a separate resource plan or RTO process, and some do not account for it yet.
- **Characteristics carried per project:** location, size, and ramp rate are nearly universal. Hourly or seasonal profile, power factor, load factor, and flexibility or curtailability are carried less consistently.

Validation, change, and RTO alignment

- **Actual vs. forecast tracking** is uneven. Several maintain customer-level variance tracking; several have no formal process yet.
- **Recent methodology change:** about half have changed approach in the last two to three years, moving toward scenario-based forecasting, milestone and certainty tracking, AMI-informed insights, quantile regression, and formalized, less subjective inclusion rules. The rest report no methodology change.
- **ISO / RTO reconciliation:** PJM utilities coordinate closely, with at least one reporting its submittal is adopted with little change; one ISO-NE utility has its large-load forecast included by the ISO; several non-RTO utilities mark this not applicable; one notes its RTO is only now beginning to produce a large-load forecast. Formal reconciliation is rare.

Biggest challenges, in the cohort's own words

The recurring themes, anonymized:

- "Uncertainty over the timeline to connect and ramp."
- "Determining the materialization of the load."
- "Filtering out customers that are fishing around for load versus ones that are legitimate."
- "The law of large numbers doesn't apply, so it is difficult to statistically quantify risk."
- "Speed. We have a total backlog of tens of gigawatts and rising monthly."
- "Executive influence on forecast inputs," when leadership wants a news-driven project reflected before it is analytically defensible.
- The subjective, non-standardized way step loads historically entered the tracker.

Where the cohort wants research

This is the clearest signal for the fall, and it lines up tightly:

- **Load materialization and realization rates** — what share of announced or contracted capacity actually shows up, and when. Multiple respondents pointed to the public data on data-center projects that never break ground.
- **Ramp characteristics and timelines** — planned versus actual ramp, by load type.
- **Load profiles and shapes by large-load type**, including seasonality, temperature sensitivity, outage frequency, and utilization or load factor.
- **Avoiding national double-counting** of the same prospective customers shopping multiple territories.
- **Industry-standard, widely accepted guidelines** for whether and how to include prospective customers in a forecast.

- **Flexible and curtailable load** treatment in the forecast.
- **AI-based forecasting methods** and better integration of interconnection-queue data.
- **Coupling with generation and gas-infrastructure adequacy** to understand the ceiling on this growth.

What this points to

The cohort is solving the same problem in parallel, with private spreadsheets and very different definitions, gates, and ramp assumptions. The strongest pull is toward a shared, peer-validated way to judge how real a prospective load is and how it will ramp, plus comparable load profiles by type. That is exactly the kind of work this group is positioned to build together, and it sets a clear agenda for the fall in-person sessions.