Highjoule

C&I Battery Storage Sizing: Check Transformer and Grid Limits Before You Pick a Cabinet

Work through C&I battery storage sizing in the right order: rank the duties, replace the annual bill with interval load data, convert delivered kWh into nameplate capacity, and check transformer headroom, breaker ratings and export limits before you choose a cabinet or container.

A single monthly or annual electricity bill, plus a transformer nameplate, gives you two numbers that bound the problem: how much energy you bought and how much power the service can carry. Neither of them tells you what size battery to buy. The bill hides the shape of your day, and the nameplate only says where the ceiling is. Owners who skip that gap usually end up asking for a size long before anyone has looked at the distribution room, and the answer they get back is a kWh figure that the transformer, the main breaker or the export rule may not allow. To avoid that, run five steps in order and stop at the fifth until it passes.

The sequence is fixed. Deciding which duty wins among several that compete for the same battery comes first. Interval load data, not annual energy, comes next, and it is what lets you state the battery output and the delivered energy the system must actually produce. That delivered figure then converts into nameplate capacity through the terms that inflate it. The result is a requirement, not yet a product: it goes back to the transformer, main breaker, busbar and point of interconnection to see whether the site can host it. Cabinet versus container is a question about the physical envelope, so it belongs at the end, once the electrical answer survives contact with the site.

Rank the duty before you rank the equipment

Commercial and industrial storage is sold as though it has one job. In practice a site usually names three or four, and they compete for the same state of charge. Over a typical project, storage can shift purchases from high-price to low-price periods, flatten the power peaks that drive demand charges, absorb surplus on-site solar generation, and shorten the exposure of critical loads to an outage. Each duty pulls the design in a different direction, and the equipment that serves one well is not automatically the equipment that serves another.

Demand-charge control is a power problem before it is an energy problem. A demand charge is billed on the maximum power consumed, typically once per month, so the quantity that matters is the kW the battery must supply during the worst metering interval of the billing period, not the kWh it can carry. Time-of-use shifting is the reverse: the deliverable is energy that can be moved into a cheaper window and recovered later, so the weekday price spread and the number of windows per day drive the design more than the peak does.

Self-consumption sits between the two. Its limit is usually the surplus the site can capture during the solar day and the evening load that can absorb it, so a large annual generation figure says little about the size that pays. Backup is the one duty that cannot share. If a critical panel must ride through an outage, that energy has to be reserved and kept out of the arbitrage schedule; a battery that spends its afternoon on price arbitrage is not a battery that answers at 2 a.m.

Rank the list before sizing anything. If peak shaving is first, you owe the design a kW number and a set of intervals to prove it. If a backup block is required, its energy is subtracted from what the daily duties can use, not added at the end as a reserve percentage someone remembers to set. Where the duties overlap, one discharge can serve both an energy charge and a billed demand, so model the dispatch once against the actual bill components rather than stacking separate schedules on top of each other.

Why a bill and a transformer nameplate stop short

Annual consumption describes volume, not behaviour. Two plants can buy the same number of kWh in a year with completely different storage requirements, because what sets the size is the sequence of intervals, and a yearly total has integrated all of it away.

Our own pre-installation input list is built around this gap. It covers three groups. The first is commercial: the electricity type, the basic demand or capacity price element, time-of-use periods and prices, and the plant's shutdown and production pattern. The second is load data: the past year's load, the average and maximum load power, and the transformer capacity. The third is site documentation: the primary power system diagram, the plant floor plan, the distribution-room layout, the cable-trench direction, and the reserved space. Only monthly or annual consumption cannot show a plant's 24-hour load pattern, so it cannot support a capacity calculation. These three groups decide the tariff basis you are chasing, the curve you size against, and where the equipment can physically connect.

Ask for the interval data at the same time step the utility uses for billing, which is commonly 15 or 30 minutes, and ask for at least twelve months of it. That resolution decides how the battery is judged: a peak that lasts one interval and a peak that lasts three hours look identical on a monthly maximum but require very different energy. Twelve months matters because the worst interval of the year is frequently not the one a friendly week in spring suggests, and a design that only covers the average month will under-deliver exactly when the bill peaks.

Collect the tariff alongside the curve. Record the periods, the demand-charge basis and its reset, whether the maximum is taken per month or over a rolling window, the operating calendar, and any export compensation or export prohibition. Note planned electrification as a separate future scenario rather than blending it into today's profile, so the sizing decision stays reproducible when the site changes.

Where intervals are missing, a short metering campaign is cheaper than a wrong answer. The failure mode here is not a small error: a design built on a guessed peak tends to be adopted after the cabinet is specified, when changing it means changing the order.

From interval data to kW and delivered kWh

With the curve in hand, the peak-shaving calculation is a row-by-row subtraction. For each interval, take the measured site load and subtract the grid import target you want to hold; if the result is positive, that is the output the battery must provide in that interval, and if it is negative or zero, the battery does nothing. The largest positive difference is your first estimate of the continuous discharge power. The energy is the sum of each interval's output times the interval length, which is simply integrating the part of the load that sits above the target.

The distinction between the two numbers is where most first estimates go wrong. Suppose four consecutive 30-minute intervals, spanning two hours, read 650, 800, 750 and 650 kW against a target of 600 kW. The required battery output is 50, 200, 150 and 50 kW. Holding that 200 kW peak across the two hours would suggest 400 kWh; integrating the actual profile gives 25 + 100 + 75 + 25 = 225 kWh delivered, with a peak requirement of 200 kW. That hypothetical profile is chosen to make the arithmetic visible, and it is not a measured site.

225 kWh is the energy that has to come out of the battery's terminals, so measure the gap in those terms from this point on. A battery sized at 225 kWh nameplate would fail this load, because the nameplate figure is not the deliverable.

One historical window is not enough either. Repeat the subtraction across the full twelve months and find the intervals that bound the design, then check the schedule that follows. After the worst hour, the battery has to recover before the next one. A cheap charging window does not by itself guarantee enough connection capacity to recharge at the rate the next peak requires. A design that discharges further than it can recover will look correct on the first day and short on the second.

From delivered energy to a nameplate

Three terms separate the delivered figure from the sticker. The usable window is the share of rated energy the control strategy is allowed to move, since a battery does not run from 0 to 100% in service. The discharge-path efficiency is the loss between stored energy and AC energy at the terminals, which includes the power conversion stage and auxiliary consumption. Capacity fade is the reduction in available energy over cycles and time, which matters for any obligation that must still be met in the later years of the project rather than on commissioning day.

The definitions here are public and worth using precisely. Round-trip efficiency is the electrical energy discharged as a percentage of the electrical energy used to charge, and it depends on operating conditions such as power level, ambient temperature and state-of-charge range. Depth of discharge is discharged energy as a percentage of rated

energy. Duration is rated energy divided by rated power. Cycle and calendar life end when available energy falls to a stated fraction of rated energy. Those definitions are the vocabulary of datasheets and standards documents, and they are what make two quotes comparable.

Read a quote with the multipliers in the open. With the 225 kWh example, a 0.90 usable window and a 0.95 discharge-path efficiency give roughly 263 kWh of initial nameplate capacity. Carrying the same obligation through to a later year at 0.80 of initial capacity gives roughly 329 kWh. Those three factors are planning assumptions chosen here to make the method visible. They are not our product parameters and they are not guarantees, and they belong in the working notes rather than in a specification.

Two errors are common at this stage. Applying the usable-window correction to a datasheet that already states guaranteed usable AC energy counts the same loss twice. And treating round-trip efficiency as a substitute for discharge-path efficiency mixes a charge-and-discharge figure with a discharge-only figure, which quietly changes the answer.

Energy also has to survive the power test. A cabinet that carries enough kWh at the wrong kW is the wrong cabinet, so the required output and the required delivered energy have to pass independently before any model is shortlisted. Power and energy are separate constraints, and neither one substitutes for the other.

Put the number back on the site node

This is the step that decides whether the project happens as planned. The battery is a new source and a new load on an existing electrical node, and that node has limits written into hardware, not in the datasheet.

Start with transformer headroom. Charging is the direction people forget: the battery charges when the plant is already drawing, so the two add. Our published guidance for low-voltage grid-connected projects is a rule of thumb: storage charging power plus the maximum load in the same period should stay below 80% of transformer capacity, so the transformer is not overloaded while the battery charges. Applied to a 1,000 kVA transformer and a 125 kW charging rate, that leaves 675 kW of coincident load headroom. If the charging-period load on that transformer exceeds it, the check fails regardless of how attractive the arbitrage looked.

The same guidance carries a second condition, for discharge: the load during the daytime peak-price period should exceed the storage discharge peak power. If the site's own load in that window is smaller than the battery's output, the surplus has nowhere to go except the grid, which turns a peak-shaving design into an export question.

Both figures are a design reference for this class of low-voltage project, and we present them as guidance rather than as a standard. They encode a margin; the transformer's own loading limits and the local code still govern the installation.

Then check the conductors and the switchgear the power passes through. The main breaker, the busbar rating and the switchgear assembly have to carry the battery's contribution on top of the existing load in the same direction, and a transformer that has headroom does not imply that a breaker or a bus section does. Reactive-power obligations belong here too, because apparent power, not active power alone, can set the real limit.

Then the export rule. Where reverse power is restricted or export is prohibited, the battery's output has to be absorbed by on-site load at every interval, not only on average. Interconnection is a regulated process, and its application content, grid-behaviour requirements and technical review are set locally; some jurisdictions allow a system to interconnect on condition that discharging back to the grid is limited. Which regime applies to your site, and how long its review takes, depends on the utility, the tariff class and the queue, so plan the schedule around your own utility's process rather than a generic estimate.

Finally, check what coincides at the point of interconnection. If rooftop or ground-mount solar is already installed, the question is not how many kW of PV and how many kW of battery exist, but which of PV, battery and grid can push power toward the same node at the same moment. An export limit that was comfortable for the PV alone may be violated by PV plus a discharging battery on a bright, low-load afternoon.

One transformer or several

In a low-voltage grid-connected project with a single transformer, the load data you provide is that transformer's load, so installed capacity can be preliminarily determined from the total load data and the transformer capacity. With several transformers running at once, the data you provide is a total across them, and a total cannot show which transformer is carrying what. Each unit has its own loading, its own headroom and possibly its own tariff basis, so per-transformer load data is what determines where the battery is installed and how large it can be at that point. Aggregating the site into one curve is convenient for the first estimate and wrong for the final one.

If a site check fails, the repair is usually one of three moves, and all three cost less before the order than after it. Reduce the kW and accept a lower peak reduction, restack the schedule so charging and discharging avoid the constrained interval, or plan a service upgrade as part of the project.

Cabinet or container comes last

Once the electrical requirement and the node checks agree, the remaining question is the envelope. Among our outdoor C&I cabinet options listed for these sites are a liquid-cooled cabinet rated 125 kW and 261 kWh, and units at 50 kW/112 kWh, 120 kW/241 kWh, 215 kW/418 kWh and 50 kW/209 kWh. Each rating is a nameplate, so the delivered-energy arithmetic above still decides how many units a peak needs, and the 125 kW tells you the power ceiling of a single unit rather than what it delivers at every condition.

Compact outdoor cabinets suit factory yards, industrial parks and buildings where floor area and cable runs are tight, and the siting inputs from the first list, the distribution-room layout, cable-trench direction and reserved space, determine which transformer the cabinet can reach and where it can stand. Containerised systems use standardised interfaces to simplify on-site wiring and arrive pre-commissioned in modular form, which is what makes them the usual answer once the energy requirement leaves single-cabinet territory and the site has the pad and access for it. The full C&I range, including container systems, is set out on our industrial and commercial energy storage solution page, and the cabinet family is listed on the C&I energy storage cabinet page.

What to have in hand before you ask for a size

If you are the owner or facility manager, the useful preparation is commercial and operational. Bring the full year of interval data at the billing interval, the tariff schedule with its demand-charge terms and reset, the production calendar and shutdown pattern, and a short list of critical loads with the ride-through you actually need. If you are the EPC or installer doing the pre-assessment, add the electrical and physical record: the single-line diagram, the plant layout, the distribution-room layout, the cable-trench route, reserved space, the transformer nameplate and count, and the main-breaker rating.

With that file, the conversation is about how much transformer headroom is available and which intervals drive the design, and it can end in a kW and kWh figure the site can actually host. Without it, a supplier can only answer with a range, and a range has no order behind it.

Our sizing inputs and the two transformer-relationship checks above sit on the C&I energy storage systems page, so you can match your file against the list before sending it. When the interval data, single-line diagram and transformer details are together, we can run the node checks against your numbers; our contact page is the starting point.

Contact Us

For catalog requests, pricing, or partnerships, please visit: https://www.hijoule.com

