RADIAN Research 100% Employee Owned

APPLICATION NOTE · PULSECONNEX

# Limiting Data Center Peak Demand with PulseConnex

How to use real-time text and email peak demand notifications to shed the right loads, in the right order, before a short peak sets the demand charge for the whole month.

## Who This Note Is For

Facility and energy managers at enterprise, colocation, and edge data centers that are billed on a demand rate, and the utility account managers who support them. It assumes the site has a utility revenue meter with a KYZ pulse output, or can have one enabled.

## Why Peak Demand Matters in a Data Center

Most commercial and industrial rates bill two things: the energy used (kWh) and the highest average demand (kW) recorded in any single demand interval during the billing period, usually 15 or 30 minutes. One bad interval sets the demand charge for the whole month. Some rates also include a ratchet, where a share of the highest peak in the past year sets a minimum billed demand for months to come, and some add time-of-use or coincident-peak charges. Check the site's tariff for the interval length, on-peak hours, and any ratchet.

A data center's IT load is high and steady, so it sets a large floor that cannot be shed. The peak is usually made by what sits on top of that floor:

- Cooling on hot afternoons, as chillers, CRAC and CRAH units, and fans work harder.
- UPS batteries recharging after an outage, a transfer test, or maintenance.
- Equipment restarting together after maintenance or a power event.
- Large scheduled compute jobs such as backups, analytics, or batch processing.
- Building loads: office HVAC, lighting, and EV chargers.

Because the floor is fixed, a shed of a few percent of total load can be a large share of the controllable load, and it only has to last long enough to keep each interval average below the target.

## What PulseConnex Adds

PulseConnex counts the same KYZ pulses the revenue meter produces, so it shows demand the way the meter measures it, updated every 5 seconds. Staff see instantaneous, average, and peak demand on a live graph from any browser or phone, and get scheduled CSV or PDF reports of 15-minute interval data for review and bill checks.

Most important for peak control, PulseConnex sends real-time peak demand notifications by text message and email. When demand crosses a limit the site sets, the people who can act are told within seconds, on their phones, wherever they are, while there is still time left in the interval to bring the average down. No one has to be watching a screen.

PulseConnex measures and alerts; it does not switch loads. The shedding is done by people following a written plan, or by the site's building management system (BMS) or a demand controller. Where automatic control is wanted, an SSI Pulse Isolation Relay lets the same KYZ output feed both PulseConnex and the control system.

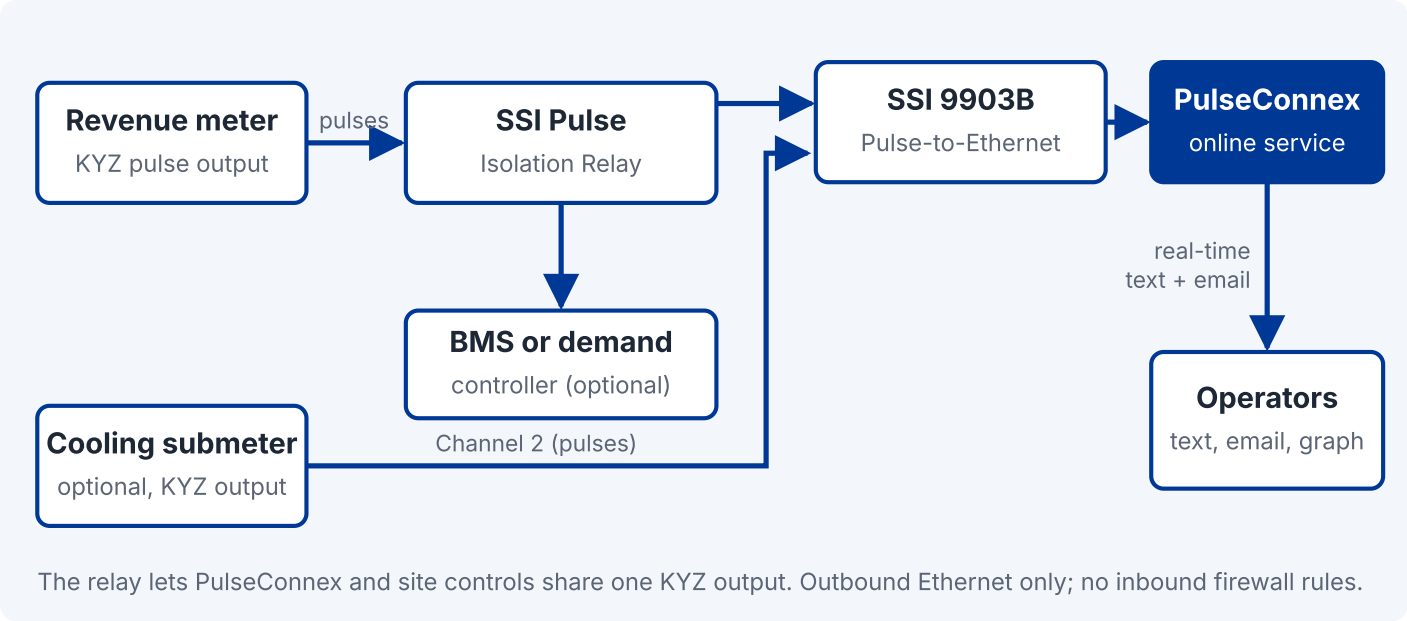



Figure 1. Connecting PulseConnex at a data center.

## Step 1: Connect

1. Ask the utility to enable a KYZ output on the revenue meter, or confirm the existing one and its pulse value (Ke, in kWh per pulse). If the output already feeds another system, add an SSI Pulse Isolation Relay so both can share it.
2. Run a two-wire control cable from the KYZ terminals to Channel 1 of the SSI 9903B Pulse-to-Ethernet Converter. A Form C (three-wire) output is wired as Form A using K and Y; enter double the meter's Form C pulse value.
3. Optional: connect a submeter with a KYZ output, such as one on the cooling plant or UPS input, to Channel 2. Seeing cooling and total load side by side shows how much of each peak is controllable.
4. Plug the 9903B into a network jack with outbound internet access and a 120 VAC outlet. It sends data out on port 80 and accepts no inbound connections.

Choose a pulse value that gives about 2 pulses per second at maximum demand and at least one pulse every 4 seconds at minimum demand. The 9903B counts up to 10 pulses per second per input.

## Step 2: Set a Demand Target and Alert Limits

1. Pull the last 12 months of bills and note the billed demand each month. Pick a target a realistic step below the typical summer peak, not below the IT floor.
2. Watch two to four weeks of PulseConnex data before acting. Find when peaks occur and what drives them; Channel 2 helps here.
3. Set a warning limit a little below the target (for example 90%) and an action limit closer to it (for example 95%). The warning starts Tier 1; the action limit starts Tier 2.
4. Set up text and email notifications for each limit. Send texts to the on-call operator and shift lead, who can act at once, and email to the facility and energy managers for the record.
5. Test the notifications at each hour of coverage: nights, weekends, and holidays included.

Remember that the bill uses the interval average, not the instantaneous reading. A short spike early in an interval may not set a peak if load comes down quickly; a moderate rise that lasts the whole interval will. The average line on the live graph is the one to watch.

## Step 3: Write the Shed Plan

List the actions in advance, in order, with an owner for each, and get them approved through the site's change process. Never include critical IT load.

### Tier 0: Prevent

Move peak-making events out of on-peak hours: UPS recharge after tests, generator transfer tests, maintenance restarts, and large batch jobs. Stagger chiller and CRAC starts. Pre-cool before on-peak hours where thermal mass or thermal storage allows.

### Tier 1: No impact

Turn off or dim lighting in unoccupied areas, set back office HVAC, pause EV charging, and defer non-urgent compute such as backups, re-indexing, and analytics.

### Tier 2: Small impact

Raise supply air or chilled water setpoints in steps while keeping server inlet temperatures inside the ASHRAE recommended range of 18 to 27 °C (64 to 81 °F). Going above that range raises server fan power and can erase the savings. Stage down redundant cooling units only where the site's design allows.

### Tier 3: Planned resources

Discharge on-site battery storage, or run generation that is permitted for non-emergency use. Emergency standby generators generally cannot be used for peak shaving: federal rules limit them to 100 hours a year of non-emergency operation, and only for purposes such as maintenance and testing. Confirm with the site's air permit.

## Step 4: Act on Real-Time Text and Email Notifications

The notification is the trigger for the shed plan. Tell the on-call team which tier each limit starts, so a text can be acted on at once.

- Warning text or email: open the live graph on a phone or computer, check the average and the time left in the interval, and start Tier 1.
- Action text or email: start Tier 2 and confirm with the shift lead. Use Tier 3 only as planned.
- Release loads in reverse order once the interval closes and demand is back below the warning limit, and watch the next interval.
- If the site has a BMS or demand controller fed from the same KYZ output, program the same tiers into it and use PulseConnex to confirm what it did.

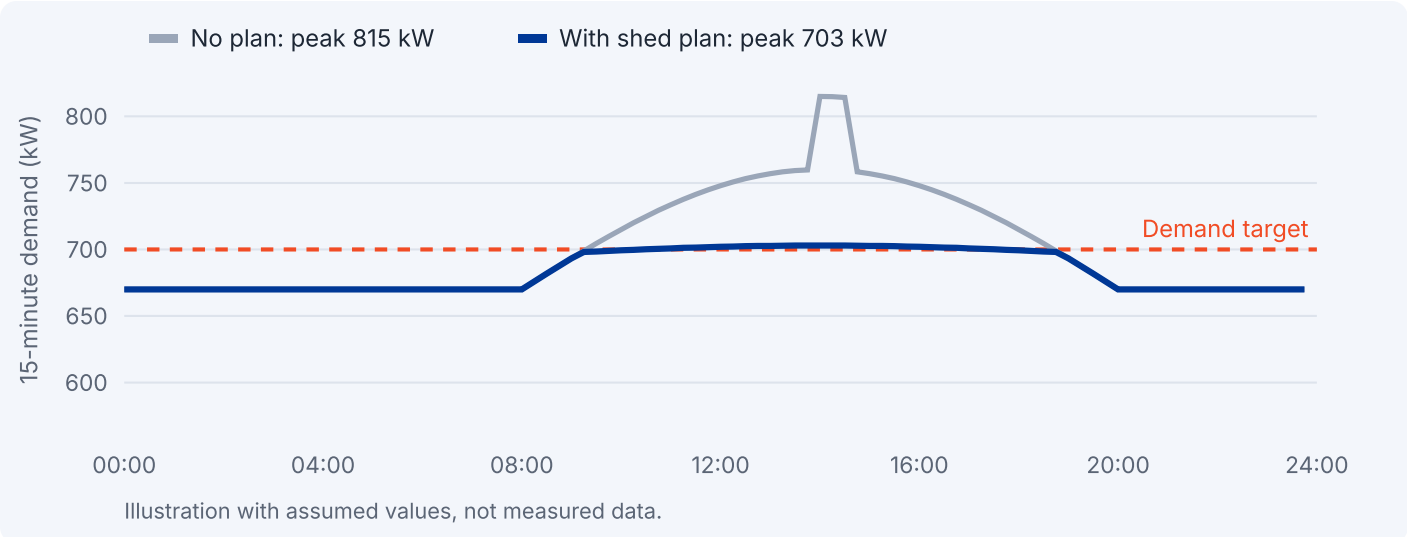



Figure 2. One summer day, with and without a shed plan (illustration).

## Step 5: Verify Results

- Schedule a monthly CSV report and compare each month's highest 15-minute interval with the target and with the billed demand.
- Review each near-miss in the load profile history: what drove it, and whether the plan acted in time.
- Adjust the target and limits each season. Cooling-driven peaks move with the weather.

## Worked Example

Assumptions, for illustration only: a site with a steady IT load near 600 kW, cooling that climbs on summer afternoons, a UPS recharge after a 2 p.m. test, a demand charge of $15 per kW per month, and a 15-minute interval. Without a plan the month's peak reaches about 815 kW. Moving the recharge to night (Tier 0) and holding the afternoon with Tier 1 and 2 actions keeps the peak near 703 kW.

That is about 112 kW lower, or about $1,680 on that month's demand charge at the assumed rate. On a rate with a ratchet, the lower peak can also lower the minimum billed demand in later months. Use the site's own tariff and data for real figures.

## Checklist

- KYZ output enabled and pulse value confirmed with the utility
- SSI 9903B installed; optional Channel 2 submeter connected
- Two to four weeks of baseline data reviewed
- Demand target, warning limit, and action limit set
- Written, approved shed plan with owners for each tier
- Text and email peak demand notifications set up and tested with the people on call
- Monthly report scheduled and reviewed against the bill

## Products Used

PulseConnex Real-Time Demand Monitor and online service; SSI 9903B Pulse-to-Ethernet Converter; SSI Pulse Isolation Relays, to share a KYZ output with a BMS or demand controller.

**See your peak before it sets your bill.**
PulseConnex Real-Time Demand Monitor from Radian Research

(765) 449-5500
sales@radianresearch.com
radianresearch.com/products/pulseconnex