Rural School District Bus Electrification

UTAH STATE UNIVERSITY ASPIRE FOUNDATION

Summary

This case study examines school-bus electrification tradeoffs using observed operational data and fleet characteristics from a rural school district in Utah. It moves from observed operations to vehicle replacement decisions, charging-system implementation, and observed-load utility-rate analysis. The workflow first estimates trip-level energy demand from measured service patterns, weather, and vehicle-use records. Those outputs are then used to test which buses can plausibly be electrified, whether replacement is cost-effective under different policy and price assumptions, what charging infrastructure would be needed to support selected electrification pathways, and how the district's current EV charging behavior affects billing risk and managed-charging opportunities.

The analysis is intended to provide practical lessons for rural school districts evaluating electrification under long routes, variable duty cycles, cold-weather operation, constrained charging infrastructure, and limited operating budgets.

The main planning result is that electrification is not a single fleetwide yes/no decision. Some buses have duty cycles that are readily compatible with depot charging and current battery capacity, while others are constrained by winter energy demand, long service days, or limited charging opportunity. Vehicle purchase incentives and utilization also matter: high-use buses can recover more value from lower operating costs, but only when the remaining upfront cost gap is not too large.

Charging design is similarly scenario-dependent. The charging analysis indicates that coordinated use of available charging power can reduce the need for immediate large power-capacity expansion in the modeled electrification pathways. The implementation insight is that port availability and power capacity should be planned separately: buses need dependable opportunities to connect during dwell and preconditioning windows, while managed charging can determine how much power each connected bus receives at each point in time.

Operational resilience is the companion to that charging-design result. The modeling shows how charging can work when energy, power, and timing are coordinated, but implementation should preserve recovery time for missed sessions, charger downtime, communication issues, and cold-weather preconditioning. An effective control strategy should prioritize readiness for the next duty cycle, not only the lowest-cost hour to deliver energy.

Report Structure

This report is organized as a four-stage planning workflow:

1. **Energy modeling**: fit and apply an explicit, temperature-sensitive per-vehicle demand model to estimate electricity and fuel use from observed trip activity.
2. **Fleet replacement modeling**: evaluate technical feasibility and lifecycle economics for vehicle replacement pathways under scenario assumptions.
3. **Charging system modeling**: optimize charging operations and infrastructure decisions over time for selected electrification trajectories, then stress-test charger use, queuing, and managed-power tradeoffs with a heuristic simulation.
4. **Charging-system solution design considerations**: use observed EV charging load to evaluate rate-schedule exposure, billing risk, and smart-energy-management opportunities.

High-Level Findings

- **Weather-sensitive energy demand is a major driver of electric-bus energy variability**, especially in winter, and should be modeled explicitly rather than absorbed into a single average kWh/mi assumption.

- **Only a subset of the fleet is consistently electrifiable under baseline assumptions** with depot-only charging and current battery constraints, indicating that feasibility is highly route- and duty-cycle-dependent.

- **Among feasible buses, cost-effectiveness is primarily governed by utilization and vehicle purchase-incentive structure.** Higher-use buses recover BEV capex more reliably through operating-cost savings, while state and federal agency purchase incentives reduce the upfront BEV cost exposure that those savings must overcome.

- **Existing depot power capacity is modeled as sufficient under coordinated operation**, even in the high-electrification pathway, when coupled with smart charging controls to allocate power across ports and buses over time. This result also depends on reliable charger operation and adequate port access.

- **Charging controls should be designed for operational recovery, not only modeled energy delivery.** Smart charging is most useful when it maintains route readiness, supports preconditioning, and leaves enough flexibility to recover from missed charging, charger outages, and late returns.

- **Utility-rate selection and smart energy management create near-term savings opportunities** for the district's current electric-bus operation. The observed charging profile is concentrated in morning and afternoon return windows, making Schedule 6A a lower-cost immediate option and managed charging a larger opportunity if route readiness is preserved.

1 Background

This case study uses a real-world rural school district in Utah as the source context for examining school-bus electrification tradeoffs. The fleet represented in this analysis includes 60 buses total, with 12 currently battery-electric vehicles (BEVs) and 48 internal-combustion-engine vehicles (ICEVs).

The source district provides an informative study context because it combines active electric-bus operations with a predominantly ICEV bus fleet, varied route-energy demands, and constrained charging infrastructure. Existing electric-bus adoption in this setting has been strongly influenced by purchase-incentive-supported procurement, which makes the case especially useful for evaluating how policy levers, operational constraints, and costs interact.

The source district operates across a large rural service area in northeastern Utah. Its operating environment includes long travel distances, substantial elevation and temperature variation, and duty cycles that can extend well beyond a typical urban school-bus schedule. The climate context is demanding for bus electrification: daily summaries from the source weather station show January averages of about 29°F high and 10°F low, August averages of about 87°F high and 55°F low, and observed historical temperatures ranging from 104°F to −17°F. Within the measured trip data, the most demanding vehicle-days exceed 500 miles and can span more than 17 hours from first departure to final return, reinforcing that the case represents a large rural service environment rather than a compact urban duty cycle.

This case study addresses four decision questions:

- Which vehicles can meet observed service requirements under current and expanded charging assumptions?
- Under consistent assumptions, how do lifecycle costs compare across Fully ICEV, Fully BEV, and cost-optimal mixed pathways?
- How sensitive are results to fuel/electricity prices, vehicle purchase-incentive levels, and BEV performance assumptions?
- What charging-system trajectory and utility-rate strategy should be carried into higher-fidelity implementation planning?

The sections that follow document the methods used to answer these questions and then summarize planning-relevant findings for district and state stakeholders, with emphasis on lessons that can inform rural school-bus electrification planning more broadly.

2 Approach

This chapter is organized as four linked analytical sections, sequenced as a planning pipeline:

1. Energy modeling (Section 3) converts operational records into trip-level electricity and fuel demand estimates using an explicit intensity model with weather and duration sensitivity.
2. Fleet composition modeling (Section 4) uses those energy and activity outputs to evaluate replacement feasibility, lifecycle cost tradeoffs, and scenario-dependent fleet-transition pathways.
3. Charging system tradeoff modeling (Section 5) takes selected fleet trajectories and evaluates both optimized coordinated charging and heuristic rule-based charger use.
4. Charging system solution design considerations (Section 6) uses observed EV charging load to evaluate rate selection, billing risk, and smart energy management for the existing charging operation.

Taken together, the structure moves from demand characterization, to fleet decision screening, to infrastructure implementation and observed-load tariff analysis. Method sections preserve technical traceability, while observations and recommendations translate the results into planning implications.

2.1 Data Description

This case study uses district fleet and operations records, including telematics records, that describe how vehicles are used today and how much energy they consume under real conditions.

The dataset represents 60 school buses, with trip activity covering October 8, 2024 through March 31, 2026. During that period, the records capture more than 90,000 individual trips, including when vehicles were in service, how far they traveled, and where trips started and ended.

To connect activity with operating cost, the analysis also uses daily fuel and electricity records from April 7, 2025 through April 4, 2026. These records include diesel, propane, and electricity use along with daily mileage, which allows the report to compare energy demand across vehicle types and seasons. A monthly electric-bus energy summary (April 2025 through March 2026) is used as an additional check on electric-bus consumption patterns. The final load-profile section also uses Fleetbox, the district's charger and telematics data platform, to evaluate observed charging behavior, rate exposure, and smart energy management opportunities from hourly energy records for May 1, 2025 through April 30, 2026 and 15-minute power records for the February-April 2026 overlap period.

In plain terms, the data provide five complementary views of district operations:

- what vehicles the district owns,
- how those vehicles are actually driven day to day,
- how much fuel or electricity that driving requires,
- how energy use changes with weather and season,
- and how the existing charging system creates hourly and sub-hourly load on the utility meter.

Together, these records create a planning baseline for evaluating where additional electrification is operationally feasible, where it is most cost-effective, and what level of charging support would be needed.

3 Energy Modeling

The objective of this step is to convert observed school-bus activity into trip-level electricity and fuel quantities. Each trip is evaluated using its distance, duration, ambient temperature, and the fitted school-bus energy/fuel coefficients. This creates a consistent estimate of BEV electricity demand and ICEV fuel demand for the same service pattern.

For observed trip i, define heating and cooling degree terms around a neutral 70°F reference temperature:

$$T_i^{\mathrm{heat}} = \max(0, 70 - T_i)$$

$$T_i^{\mathrm{cool}} = \max(0, T_i - 70)$$

The modeled BEV propulsion energy, temperature-sensitive energy term, and total trip energy are:

$$E_i^{\mathrm{prop}} = d_i e_{\mathrm{flat}}$$

$$E_i^{\mathrm{temp}} = h_i\big(e_{\mathrm{heat}}\, T_i^{\mathrm{heat}} + e_{\mathrm{cool}}\, T_i^{\mathrm{cool}} + e_{\mathrm{aux}}\big)$$

$$E_i^{\mathrm{trip}} = E_i^{\mathrm{prop}} + E_i^{\mathrm{temp}}$$

The modeled ICEV fuel equation uses the same structure:

$$F_i^{\mathrm{prop}} = d_i f_{\mathrm{flat}}$$

$$F_i^{\mathrm{temp}} = h_i\big(f_{\mathrm{heat}}\, T_i^{\mathrm{heat}} + f_{\mathrm{cool}}\, T_i^{\mathrm{cool}} + f_{\mathrm{aux}}\big)$$

$$F_i^{\mathrm{trip}} = F_i^{\mathrm{prop}} + F_i^{\mathrm{temp}}$$

where:

- i: observed trip.
- E_i^{prop}: modeled BEV propulsion energy for trip i.
- E_i^{temp}: modeled BEV temperature-sensitive energy term for trip i; the fitted term may include cabin HVAC, battery conditioning, battery-performance impacts, and other temperature-related effects.
- E_i^{trip}: total modeled BEV trip energy.

Table 1: Trained energy-model coefficients used in the district dataset.

e_{flat} (kWh/mi)	e_{heat} (kWh/hr/°F)	e_{cool} (kWh/hr/°F)	e_{aux} (kWh/hr)	f_{flat} (gal/mi)	f_{heat} (gal/hr/°F)	f_{cool} (gal/hr/°F)	f_{aux} (gal/hr)
1.1482	0.2647	0.1635	0.0000	0.0708	0.0000	0.0110	2.7634

- d_i: trip distance in miles.
- h_i: trip duration in hours.
- T_i: ambient temperature in °F.
- T_i^{heat} and T_i^{cool}: heating and cooling degree terms relative to the 70°F reference point.
- $e_{\cdot}$: electricity model coefficients (kWh terms).
- $f_{\cdot}$: fuel model coefficients (gallon terms).

3.1 Parameter Training and Application

The coefficients above were trained from observed trip-level energy/fuel rows in the data:

1. Filter to observed rows with positive distance and duration.
2. Build a design matrix using:
   - distance, for propulsion energy or fuel demand;
   - trip duration multiplied by T_i^{heat}, for heating-sensitive demand;
   - trip duration multiplied by T_i^{cool}, for cooling-sensitive demand;
   - trip duration, for auxiliary time-based demand.
3. Fit weighted least squares (trip miles as weights) separately for electricity and fuel.
4. Apply non-negativity constraints so propulsion intensity remains positive and temperature-sensitive/auxiliary terms remain nonnegative.
5. Save the trained coefficients for reuse in dashboard scenario evaluation and downstream feasibility, cost, and charging-system calculations.

3.2 Trained Coefficients Used in District Dataset

Temperature and duration terms remain active in the fitted model and are the main source of seasonal variation represented in this analysis. The zero BEV auxiliary coefficient should not be interpreted as an absence of auxiliary electric loads. In the BEV fit, the heating- and cooling-duration terms emerge as better empirical representations of observed auxiliary and thermal loads. By contrast, the ICEV fit retains a nonzero auxiliary coefficient because idling and other temperature-independent non-propulsion fuel use appear more consistent across temperature conditions in the observed data. These coefficients are empirical planning terms, not a complete physical decomposition of every accessory load.

3.3 Observations

3.3.1 Winter Conditions Increase Electric-Bus Energy Use

Figure 1 shows the temperature-sensitive component of electric-bus energy demand directly from the fitted intensity model, isolating the temperature-driven term from propulsion load and making the modeled effect interpretable as an average energy-consumption rate. The figure uses kWh/hr, equivalent to average kW over the modeled hour. In this model form, the temperature-driven term grows as ambient temperature moves away from 70°F, and the slope on the heating side is steeper than the cooling side.

That asymmetry matters for rural school-district operations. The heating coefficient is larger than the cooling coefficient, so cold-weather days create a disproportionate energy burden compared with warm-weather days of similar service length. Because the temperature-sensitive term is time-based rather than distance-based, this burden is amplified on days with low average speed, long dwell/idle time, or short mileage spread over long service hours.

Figure 2 shows this effect at the modeled vehicle-day level: as temperatures decline, the temperature-sensitive term can account for a much larger share of daily energy, with some moderate-dis-

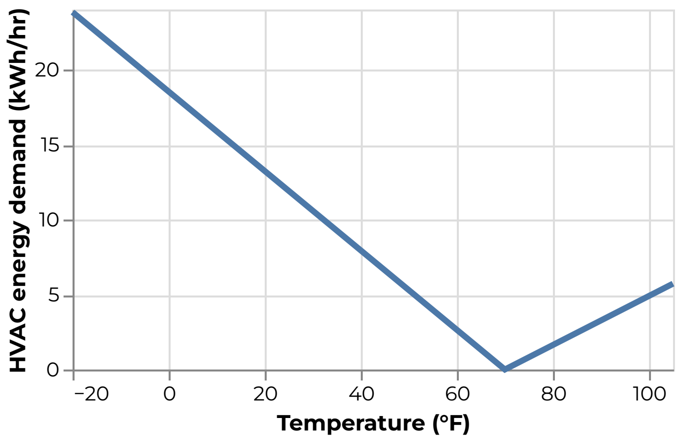

Figure 1: Modeled temperature-sensitive electric-bus energy demand rate from the fitted school-bus intensity model, shown as kWh/hr versus ambient temperature.

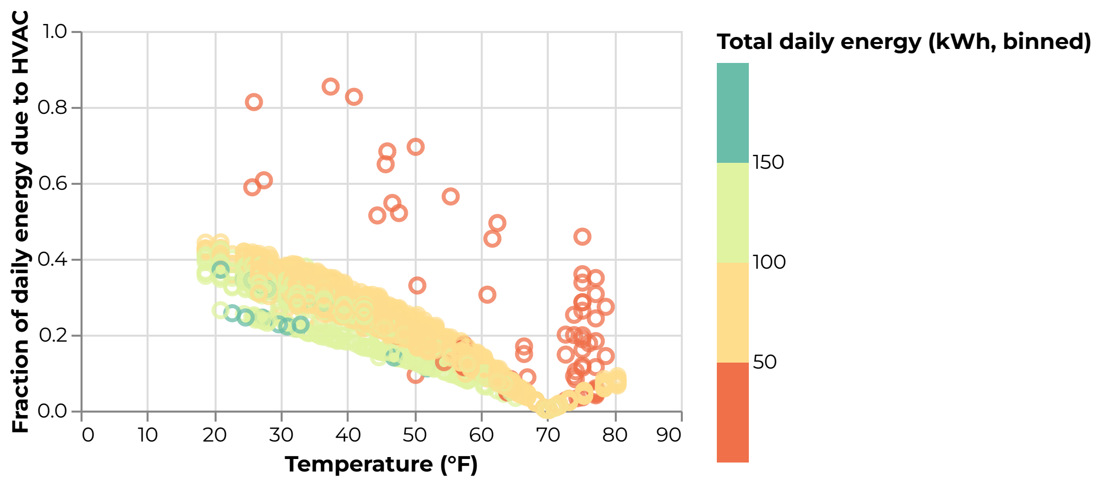

Figure 2: Modeled BEV vehicle-day temperature-sensitive energy share versus temperature, colored by total daily modeled energy (kWh).

tance cold days approaching a 50% share. This is exactly the operating regime where range-risk increases most quickly, even when miles traveled are moderate. (However, as illustrated by the color scale, days with extremely high temperature-sensitive energy fractions tend to be the lowest-overall-energy days, where range risk is not a factor.)

For planning, the temperature-sensitive result should be read as a combined cold-weather performance burden rather than as a cabin-heating estimate alone. The fitted term may reflect cabin heat, battery conditioning, reduced battery performance, and other weather-related effects that occur together in the observed data. That makes connected parking, preconditioning, and sufficient ready-time buffer especially important: they help the fleet manage the cold-weather burden before buses leave the depot, rather than discovering the energy penalty during the route.

3.3.2 Weather Explains Much, but Not All, of the Variation

Figure 3 compares measured daily behavior to values reconstructed from the fitted model. Ambient temperature is clearly a strong driver, but the vertical spread at similar temperatures indicates that additional variables still matter: route topology, stop density, speed profile, layover behavior, and driver/service pattern all contribute to day-level variation that temperature alone cannot explain.

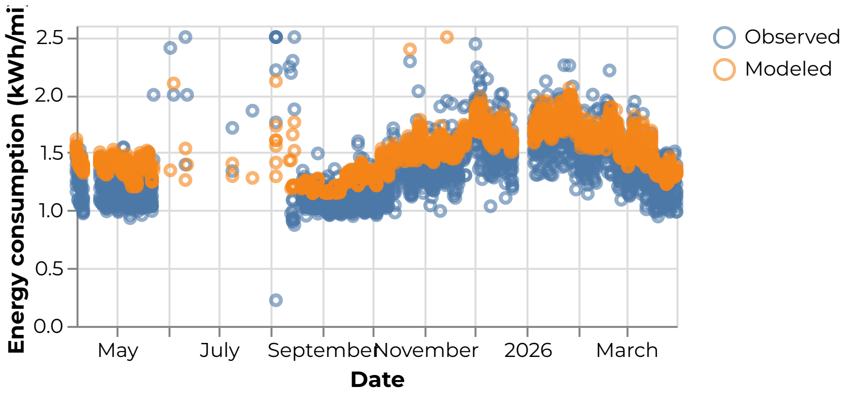

Figure 3: Comparison over time between observed daily electric-bus energy intensity and model-predicted intensity. Temperature explains the main seasonal structure, while remaining spread reflects route and usage effects outside the modeled covariates.

This spread is expected in a planning model that prioritizes broad robustness over exhaustive route-level detail. The key interpretation is not that the model captures every fluctuation, but that it captures the dominant thermal trend while preserving a realistic range of operational outcomes.

Figure 3 also suggests a conservative tendency in modeled intensity: predicted points generally sit nearer the upper envelope of comparable observed clusters than the lower envelope. That conservative tilt is intentional and feeds directly into the range-feasibility screen in Section 4 because it affects which vehicles are treated as electrifiable.

3.3.3 Conservative Assumptions Reduce Underestimation Risk

Figure 4 evaluates prediction error in kWh totals rather than only per-mile energy intensity. This framing is important for planning because charging feasibility and depot energy sizing depend on aggregate daily demand, not only per-mile efficiency.

The error pattern leans positive, meaning modeled totals are more likely to overestimate than underestimate observed energy. In this context, that is intentional: overestimation reduces the chance of falsely declaring a vehicle electrifiable when it may fail under real operating variability. In other words, the model is intentionally biased toward operational safety margins.

At the same time, the points remain concentrated near the $y = x$ diagonal, where modeled and observed totals are equal. (Perfect alignment with this diagonal would mean the model exactly predicts actual energy consumption.) This indicates that conservative bias is controlled rather than arbitrary. The model preserves strong agreement with observed totals while still protecting against optimistic underprediction in downstream feasibility screening.

4 Fleet Composition Modeling

Fleet composition modeling translates the trip-level energy estimates into vehicle replacement decisions. This stage asks two linked questions: which buses can satisfy their observed service with a BEV replacement, and among those technically feasible replacements, which are economically preferred under each scenario.

4.1 Replacement Candidates

Because this case study only includes school buses, the replacement-choice structure is simple: each vehicle can either remain or be replaced with an ICEV bus or a BEV bus, subject to scenario rules and BEV feasibility. The Fully ICEV scenario only allows ICEV operation. The Fully BEV scenario assigns BEV replacements where the range-feasibility screen is passed and retains ICEV operation where it is not. The Optimal scenario allows feasible vehicles to choose the lower-

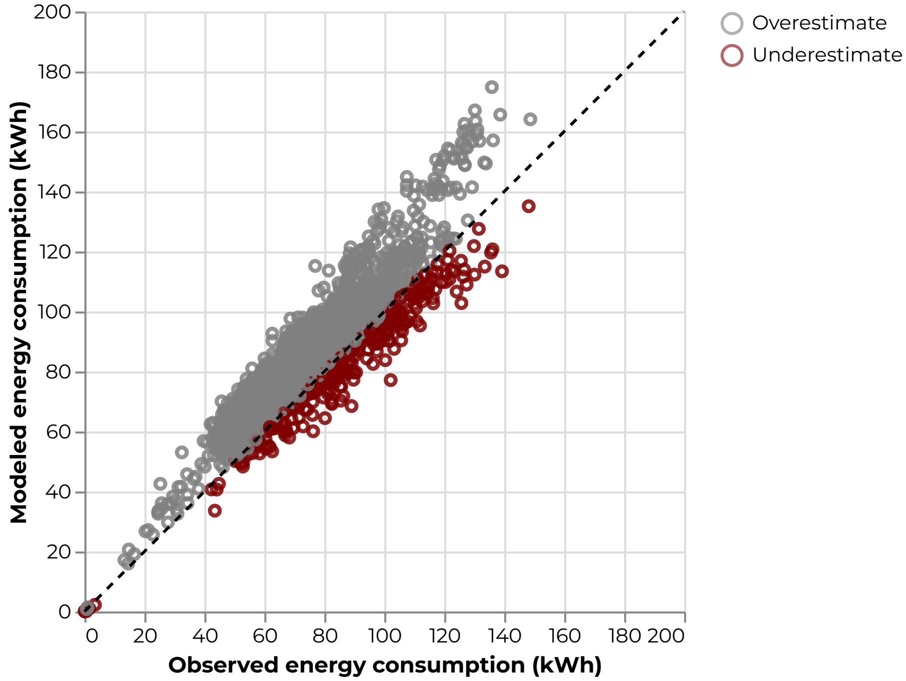

Figure 4: Comparison between energy consumption totals observed in the data and totals predicted by the model. The prediction error is intentionally biased toward overestimation of energy demand.

Table 2: Key parameters for the ICEV and BEV replacement candidates.

Total ICEV capex ($)	ICEV capex after incentive ($)	Total BEV capex ($)	BEV capex after incentive ($)	BEV usable battery (kWh, modeled constraint)	BEV nominal range (mi)
154,000	107,800	475,000	175,000	155.2	150

total-cost ICEV or BEV pathway, while infeasible vehicles remain ICEV.

Each candidate is parameterized by purchase cost and fuel/electric intensity multipliers relative to the modeled baseline demand from the energy model section. BEV feasibility is evaluated against usable battery energy in kWh. Key parameters for the replacement candidates are given in the table below:

The usable battery value is the modeled energy available for service after applying the assumed reserve/buffer strategy; it is lower than nominal pack capacity and is the value used in the feasibility constraint.

These inputs define the two fuel pathways evaluated by the scenario engine. Initial fleet mapping is shown below:

Initial ICEV count	Initial BEV count	Total
48	12	60

4.2 Range Feasibility Modeling

Range feasibility is evaluated at the vehicle-day level using modeled trip energy plus charging-opportunity structure. The trip-energy term below is the same total BEV trip-energy concept defined in Section 3, with additional indices for vehicle, day, range chain, and trip position within that chain.

For each vehicle, trips are segmented into consecutive “range chains,” meaning the portion of a vehicle duty cycle that must be completed before the next charging opportunity. A new chain begins when either:

1. the previous trip ended at a designated charging stop and the dwell time before the next trip is at least 20 minutes, or
2. a new calendar day begins.

Within each chain, cumulative BEV energy is:

$$C_{v,d,k,j} = \sum_{q=1}^{j} E_{v,d,k,q}^{\text{trip}}$$

where:

- v: vehicle index.
- d: observed service-day index.
- k: range-chain index within a vehicle-day.
- j: trip position within a range chain.
- q: summation index over trips up to trip position j.
- $E_{v,d,k,q}^{\text{trip}}$: scenario-adjusted modeled BEV energy for vehicle v, day d, range chain k, and trip position q, using the trip-energy model from Section 3.
- $C_{v,d,k,j}$: cumulative energy required from the start of range chain k through trip position j.

The required daily energy for BEV service is the maximum cumulative chain energy that day:

$$R_{v,d} = \max_{k,j} C_{v,d,k,j}$$

This represents the minimum usable battery energy needed to complete the most demanding between-charge segment on that day.

The model then compares each vehicle-day requirement $R_{v,d}$ with the usable BEV battery energy available under the scenario assumptions. A vehicle-day is treated as electrifiable when the required between-charge energy fits within usable battery capacity. A vehicle is treated as feasible when at least 95% of its observed travel days are electrifiable.

Baseline range-feasibility outcomes (60 vehicles):

Metric	Value
Feasible vehicles	30
Infeasible vehicles	30
Feasible share	50%
Infeasible share	50%

Vehicle-level electrifiable-day-share results under baseline settings:

Feasibility group	Vehicles	Mean electrifiable-day share	Median electrifiable-day share
Feasible	30	0.9913	1.0000
Infeasible	30	0.6249	0.6351
Total	60	0.8081	0.9531

Across vehicles, electrifiable-day shares range from about 6% to 100%, showing that many units are materially constrained under the base charging assumptions.

4.3 Cost-Effectiveness Modeling

Cost effectiveness is evaluated with a constrained net-present-value (NPV) replacement policy for each vehicle, then aggregated to scenario totals.

For vehicle v at decision year index t:

- $a_{v,t}$ is asset age in years,
- $f_{v,t} \in \{\text{icev}, \text{bev}\}$ is existing fuel type,
- a_v^{min} and a_v^{max} are minimum and maximum replacement ages,
- r is discount rate.

Discount factor:

$$\delta_t = (1+r)^{-t}$$

4.3.1 Replacement-Age Constraints

The policy enforces:

$$a_{v,t} < a_v^{min} \Rightarrow \text{do not replace}$$

$$a_{v,t} \geq a_v^{max} \Rightarrow \text{replace this year}$$

In this case study, the replacement-age bounds used by the model are:

Min replacement age	Max replacement age
15	20

4.3.2 Incentive and Cost Scaling

Capital cost for each replacement option is adjusted by scenario-specific market-change and vehicle purchase incentive assumptions. The market-change parameter represents the assumed change in vehicle purchase price relative to the base case. The purchase-incentive coverage parameter represents the share of eligible vehicle purchase cost assumed to be covered by state and/or federal agency funding. The resulting effective purchase cost is then used in the replacement cash-flow comparison.

4.3.3 Annual Operating-Cost Treatment

Annual operating energy is taken from the modeled demand profiles derived from the trip-level intensity model. For each vehicle and candidate fuel option, the model computes annual operating energy cost by summing the relevant energy purchases across the modeled year. ICEV operation sums modeled liquid-fuel consumption multiplied by the scenario diesel fuel price. BEV operation sums modeled charging-energy consumption multiplied by the scenario electricity price. These annual energy-cost totals then enter the vehicle cash-flow comparison alongside fixed and replacement-related costs. Later charging-system results replace the scalar BEV electricity price with explicit utility-bill calculations for selected electrification trajectories.

Total annual cash flow $\mathrm{CF}_{v,t}$ for a decision includes:

- one-time costs in replacement years (e.g., capex),
- annual fixed costs,
- per-mile costs,
- annual energy-cost total.

Discounted cash flow contribution in year t:

$$\mathrm{DCF}_{v,t} = \delta_t \cdot \mathrm{CF}_{v,t}$$

4.3.4 Terminal Value

Terminal value is included because vehicles purchased near the end of the analysis horizon may still have useful life remaining after the final modeled year. Crediting that remaining value avoids unfairly penalizing late-period replacements relative to vehicles that were purchased earlier and used for more of their life within the modeled horizon.

If a vehicle is replaced within the horizon, residual value is credited at horizon end:

$$\mathrm{TV}_v = \mathrm{Capex}^{\mathrm{eff}}_{v,f} \cdot \frac{\max(U_v - a_{v,T+1}, 0)}{U_v}$$

where U_v is useful life (set to a_v^{max} in this implementation), and the model adds $-\delta_T \cdot \mathrm{TV}_v$ as a final-year cash-flow component.

Terminal-value terms:

- T: analysis-horizon length in years (here 15).
- $a_{v,T+1}$: projected vehicle age immediately after the final modeled year.
- δ_T: terminal discount factor at horizon end.

4.3.5 Keep-vs-Replace Decision Rule

At each state (t, a, f, p) (year index, age, fuel type, whether purchased in-horizon), V_t represents the minimum discounted cost-to-go from year t through the end of the planning horizon. The decision rule is recursive: keeping a vehicle adds the same-year keep cash flow and moves the state to the next year, while replacing a vehicle adds the same-year replacement cash flow and then evaluates the next year's cost-to-go for the new fuel option.

The policy solves:

$$V_t(a,f,p) = \begin{cases} \min\left\{V_t^{\mathrm{keep}}, V_t^{\mathrm{replace}}\right\}, & a_v^{min} \le a < a_v^{max} \\ V_t^{\mathrm{keep}}, & a < a_v^{min} \\ V_t^{\mathrm{replace}}, & a \ge a_v^{max} \end{cases}$$

with:

$$V_t^{\mathrm{keep}} = \delta_t \cdot \mathrm{CF}_t^{\mathrm{keep}} + V_{t+1}(a+1, f, p)$$

$$V_t^{\mathrm{replace}} = \min_{h \in \mathcal{F}_v^{\mathrm{scenario}}} \left[\delta_t \cdot \mathrm{CF}_t^{\mathrm{replace}}(h) + V_{t+1}(1, h, 1)\right]$$

The allowed replacement-fuel set $\mathcal{F}_v^{\mathrm{scenario}}$ follows the scenario and BEV-feasibility rules defined in the preceding subsections.

Decision-state terms:

Table 3: Dashboard planning inputs and model use.

Input	How it is used
Electricity price	Converts modeled annual BEV electricity demand into annual operating cost for dashboard sensitivity analysis.
Fuel price	Converts modeled annual ICEV fuel demand into annual operating cost.
ICEV vehicle purchase-cost adjustment	Raises or lowers the assumed ICEV replacement purchase cost relative to the base case.
BEV vehicle purchase-cost adjustment	Raises or lowers the assumed BEV replacement purchase cost relative to the base case.
ICEV vehicle purchase-incentive coverage	Reduces the modeled eligible ICEV replacement purchase cost by the assumed state and/or federal agency funding share.
BEV vehicle purchase-incentive coverage	Reduces the modeled eligible BEV replacement purchase cost by the assumed state and/or federal agency funding share.
Electrifiable-day threshold	Sets the share of observed travel days that must fit within usable BEV battery capacity before a bus is treated as technically feasible.
BEV range increase	Tests improved usable range or future battery capability relative to the base modeled BEV.
Temperature-sensitive energy multiplier	Tests sensitivity to higher or lower weather-related energy demand, including heating, cooling, battery-conditioning, and other temperature-sensitive effects represented in the fitted model.
Propulsion energy multiplier	Tests sensitivity to higher or lower driving energy demand.
Ambient temperature offset	Shifts the temperature input used in the temperature-sensitive energy model.
Allowed non-depot charging stops	Tests whether additional charging opportunities outside the depot would change feasibility.
Time horizon	Sets the number of years included in lifecycle-cost comparison.
Discount rate	Converts future costs into present-value terms.
Simulation start year	Sets the first modeled year for the fleet-transition calculation.

- $p \in \{0, 1\}$: indicator for whether the vehicle has already been purchased or replaced within the modeled horizon; this state tracks in-horizon replacement status for residual-value treatment and prevents the decision rule from treating an old carried-in vehicle the same as a newly purchased one.
- h: candidate replacement fuel choice in the decision year.
- $\mathrm{CF}_t^{\mathrm{keep}}$: same-year cash flow if the vehicle is retained.
- $\mathrm{CF}_t^{\mathrm{replace}}(h)$: same-year cash flow if replaced now with fuel option h.

4.4 Decision Support Dashboard

As part of this study, a web-based decision-support dashboard was developed to help the district explore scenarios and evaluate how selected planning variables affect electrification feasibil-

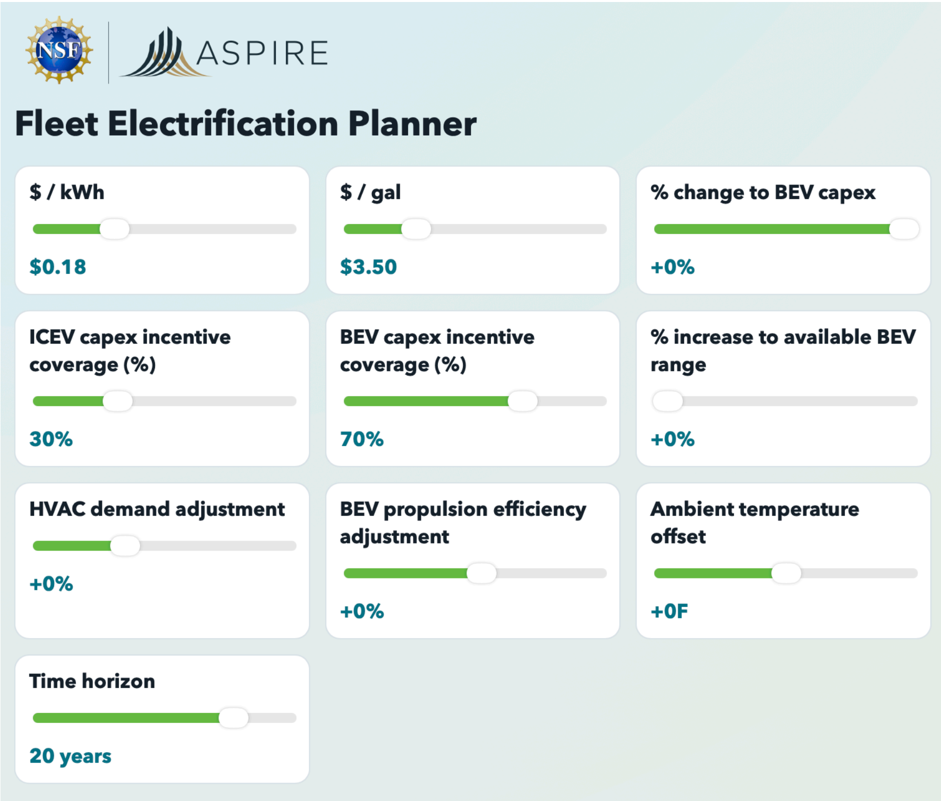

Figure 5: Dashboard slider panel. The depicted sliders are used for the planning inputs listed above.

ity, lifecycle cost, and replacement timing. For each user-selected parameter vector $\mathbf{x}$, the dashboard evaluation engine compares three scenarios ("Fully ICEV," "Fully BEV," and "Optimal") on a common fleet/activity baseline.

4.4.1 Inputs

The dashboard controls the model through reader-facing planning inputs rather than requiring the user to interact with the equations directly:

4.4.2 Compute-Response Contract

For each request, the dashboard executes:

$$\mathcal{R} = F(\mathbf{x})$$

where $F(\cdot)$ is the scenario engine and $\mathcal{R}$ contains:

- scenario NPV totals,
- annual and component-level discounted cash flows,
- vehicle-year fuel-state trajectories,
- vehicle purchase timelines,
- per-vehicle range-feasibility metrics,
- per-vehicle diagnostic statistics.

Core derived metrics shown in the dashboard include:

$$T_s = \sum_{y,c} \mathrm{DCF}_{s,y,c}$$

$$S_{\text{icev}} = T_{\text{Fully ICEV}} - T_{\text{Optimal}},$$
$$S_{\text{bev}} = T_{\text{Fully BEV}} - T_{\text{Optimal}}$$

where T_s is total discounted cost for scenario s, y indexes modeled years, and c indexes cost components. S_{icev} is the discounted-cost savings of the Optimal scenario relative to the Fully ICEV benchmark, and S_{bev} is the discounted-cost savings of the Optimal scenario relative to the Fully BEV benchmark.

4.4.3 Output Views

The following figures illustrate example dashboard visualizations. Actual dashboard outputs vary with input selections, and the same interface can be rerun as fuel costs, vehicle costs, incentive

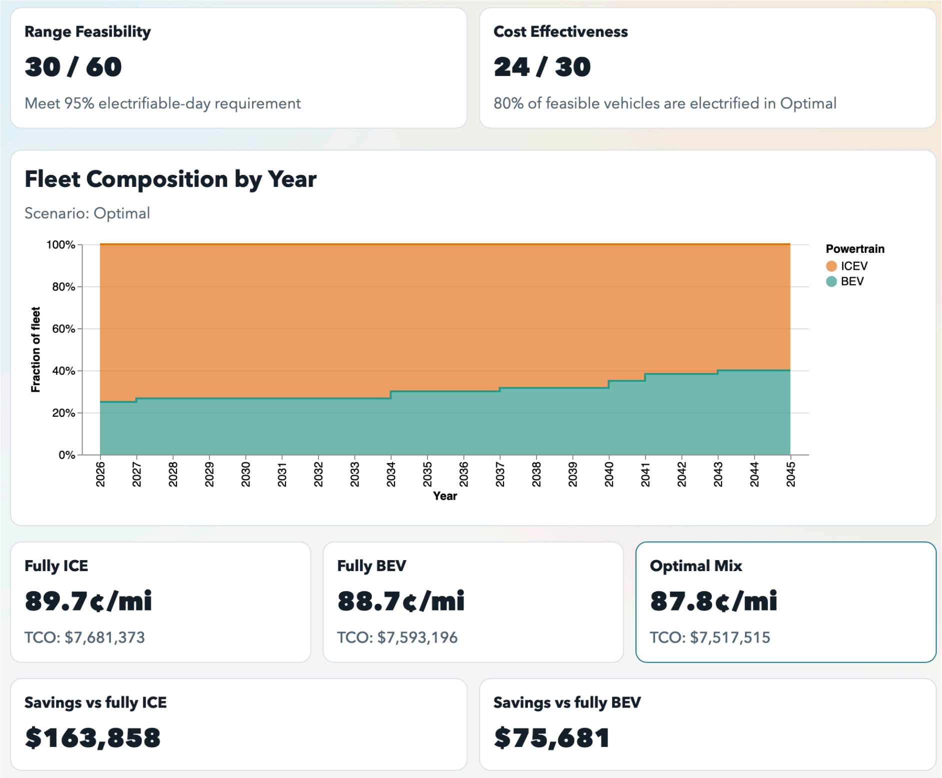

Figure 6: Example dashboard summary/feasibility cards and top-line scenario outputs.

availability, battery sizes, or operating assumptions change.

The reporting layout is card-based and includes:

- feasibility and summary metric cards (vehicle feasibility, electrification count, scenario totals, savings; Figure 6),
- stacked cost breakdown by scenario and component (Figure 7),
- annual cash-flow chart by component (Figure 7),
- fleet-composition trajectory by year and scenario (Figure 6),
- vehicle-purchase timeline by year/fuel type (Figure 8),
- vehicle-selection diagnostics (including day-level range requirement checks; Figure 6),
- movement map view for selected activity windows (Figure 9 and Figure 10),

4.4.4 Role in Decision-Making

Methodologically, the dashboard functions as a controlled sensitivity environment: each slider change updates $\mathbf{x}$, re-solves the scenario set, and returns comparable outputs under identical structural assumptions. This enables transparent tradeoff analysis across cost, feasibility, and transition timing without changing the underlying fleet/trip evidence base.

4.5 Observations

4.5.1 Half of the Fleet Is Feasible Under Baseline Battery Assumptions

Under baseline assumptions, 30 of 60 buses (50%) are feasible to electrify with the modeled school-bus battery configuration (194-kWh nominal pack, 155.2-kWh usable). The vehicle-level required-energy figure shows that a large share of buses sits close to the operational feasibility boundary, rather than separating into two clearly distinct groups.

This “near-threshold” structure is operationally important. The feasible share can increase to roughly three-quarters of the fleet under changes such as (a) increasing available on-board energy by about 30%, or (b) reducing temperature-sensitive draw from traction battery demand through more efficient vehicle thermal

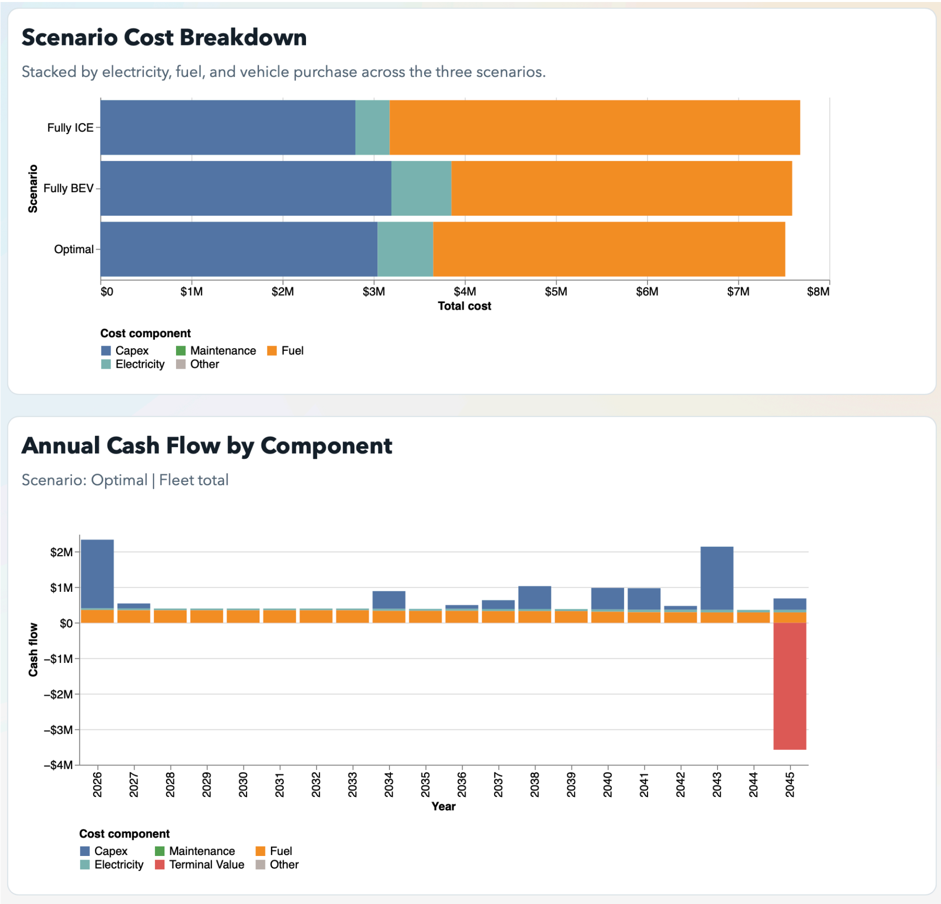

Figure 7: Example dashboard total-cost and annual cash-flow views by scenario/component.

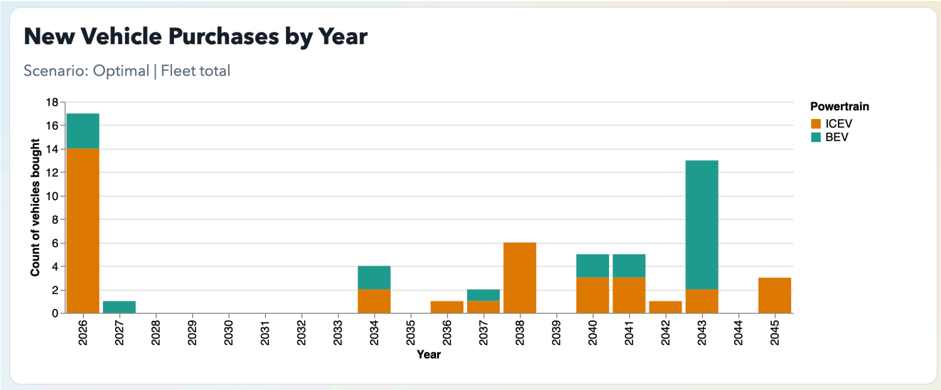

Figure 8: Example dashboard vehicle-purchase timeline by year and fuel type.

management, such as heat pumps, improved battery conditioning, or preconditioning strategies that use grid energy before departure. The remaining infeasible buses appear to be the vehicles with irregular/high-demand service patterns (e.g., athletics or field-trip style duty), where electrification likely requires more structural changes: higher-energy vehicles, aggressive mid-day fast charging, or additional en-route/dynamic charging support.

Many routes are also long enough that feasibility depends on restoring energy in between morning and afternoon service, not only overnight depot charging.

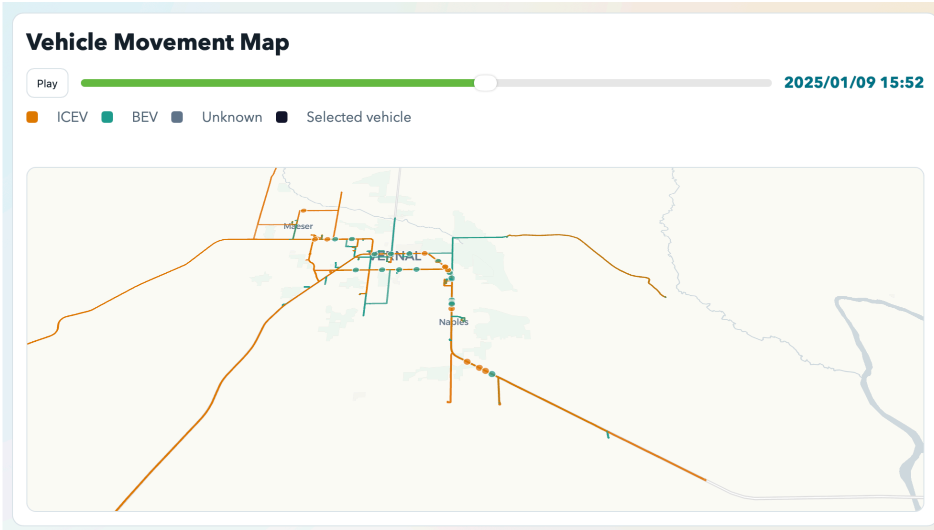

Figure 9: Example dashboard map card at compact view scale for selected activity windows.

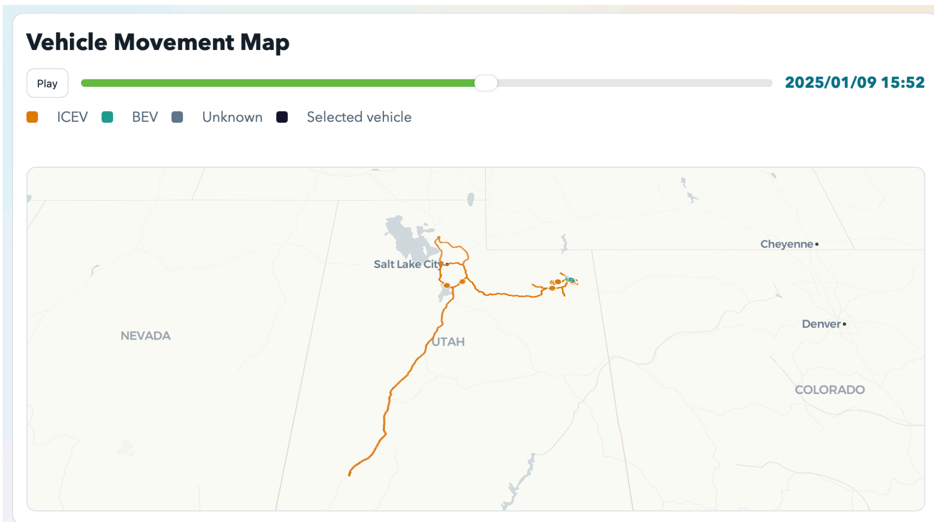

Figure 10: Example dashboard map card at expanded view scale for selected activity windows.

In the vehicle-level required-energy figure, each row summarizes the distribution of modeled daily required energy for one bus, ordered by the bus's 95th-percentile daily requirement and split into two columns for readability. Wider or farther-right distributions indicate buses with more frequent high-energy service days.

Sparse trip records should be interpreted carefully in this figure. Bus 99, for example, appears on only 12 distinct travel dates spread sparsely across the full time domain of the trip data. That pattern could represent a very sporadically used bus, or it could indicate a telematics assignment or data-quality issue. Under the baseline screen, Bus 99 is classified as infeasible because one of those 12 travel days requires too much energy for the modeled BEV case; with only 12 observed travel days, a single infeasible day is enough to reduce the electrifiable-day share below the 95% threshold.

4.5.2 Cost-Effectiveness Depends on Recovering Capex Through Operations

In this study, cost-effectiveness is driven by whether operating-cost savings can overcome BEV-ICEV capex differences over the planning horizon. The baseline vehicle purchase-incentive assumptions are aligned with Utah's clean school

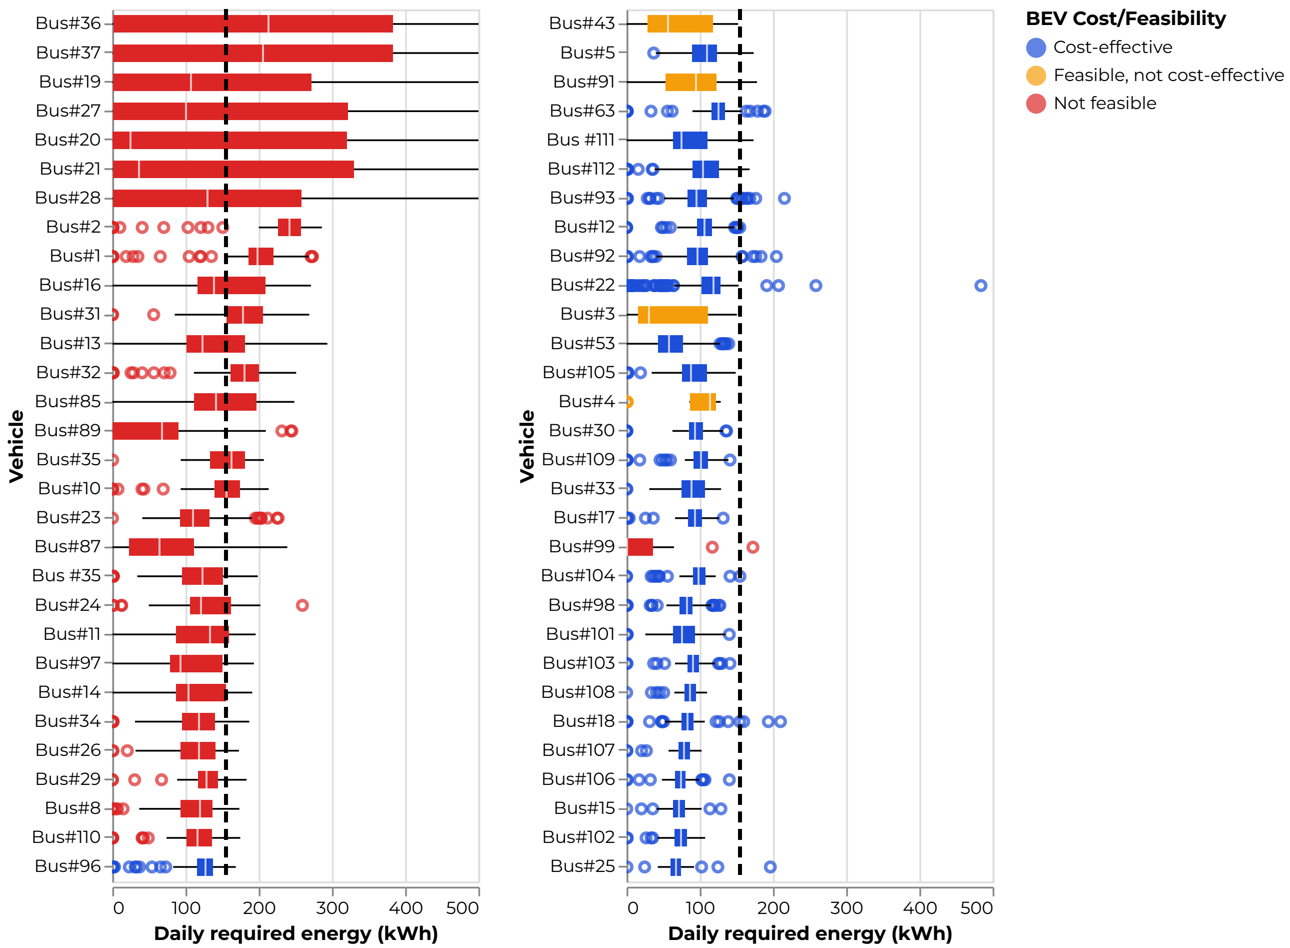

Figure 11: Vehicle-level required daily energy distributions, colored by cost/feasibility class. (Cost-effectiveness assumes baseline incentive coverage described in the next section.) Each row represents one bus. The box spans the middle 50% of modeled daily requirements, the center mark indicates the median, whiskers extend to the most extreme non-outlier values within 1.5 times the interquartile range, and circles mark outlier days. The vertical dashed line marks the modeled usable BEV battery threshold of 155.2 kWh; buses with distributions extending beyond that line are more likely to fail the baseline feasibility screen.

bus funding context. The BEV incentive assumption is based on the U.S. Environmental Protection Agency's Clean Heavy-Duty Vehicles (CHDV) Grant, School Bus Sub-Program, administered in Utah by the Utah Department of Environmental Quality (UDEQ), Division of Air Quality (UDAQ). Under the applicable funding round, eligible ADA-compliant electric school buses may receive up to $300,000, while non-ADA-compliant buses may receive up to $280,000. The model assumes the higher ADA-compliant award, equivalent to approximately 63% of the modeled $475,000 BEV purchase cost. These funds are limited to the applicable grant period and are subject to program eligibility, application, award, and vehicle-replacement requirements; they should not be interpreted as a permanent or universally available subsidy (Utah DEQ Clean School Bus Program).

The low-electrification case assumes 30% ICEV purchase-incentive coverage for eligible non-EV replacements. This is a conservative scenario assumption informed by Utah clean-fleet funding levels, including DERA-style support of up to approximately 35%; it is not intended to represent a guaranteed award or the maximum available under every program. Other funding sources, including Volkswagen Environmental Mitigation Trust support, may provide higher coverage—up to approximately 55% in some circumstances—but eligibility, funding availability, vehicle require-

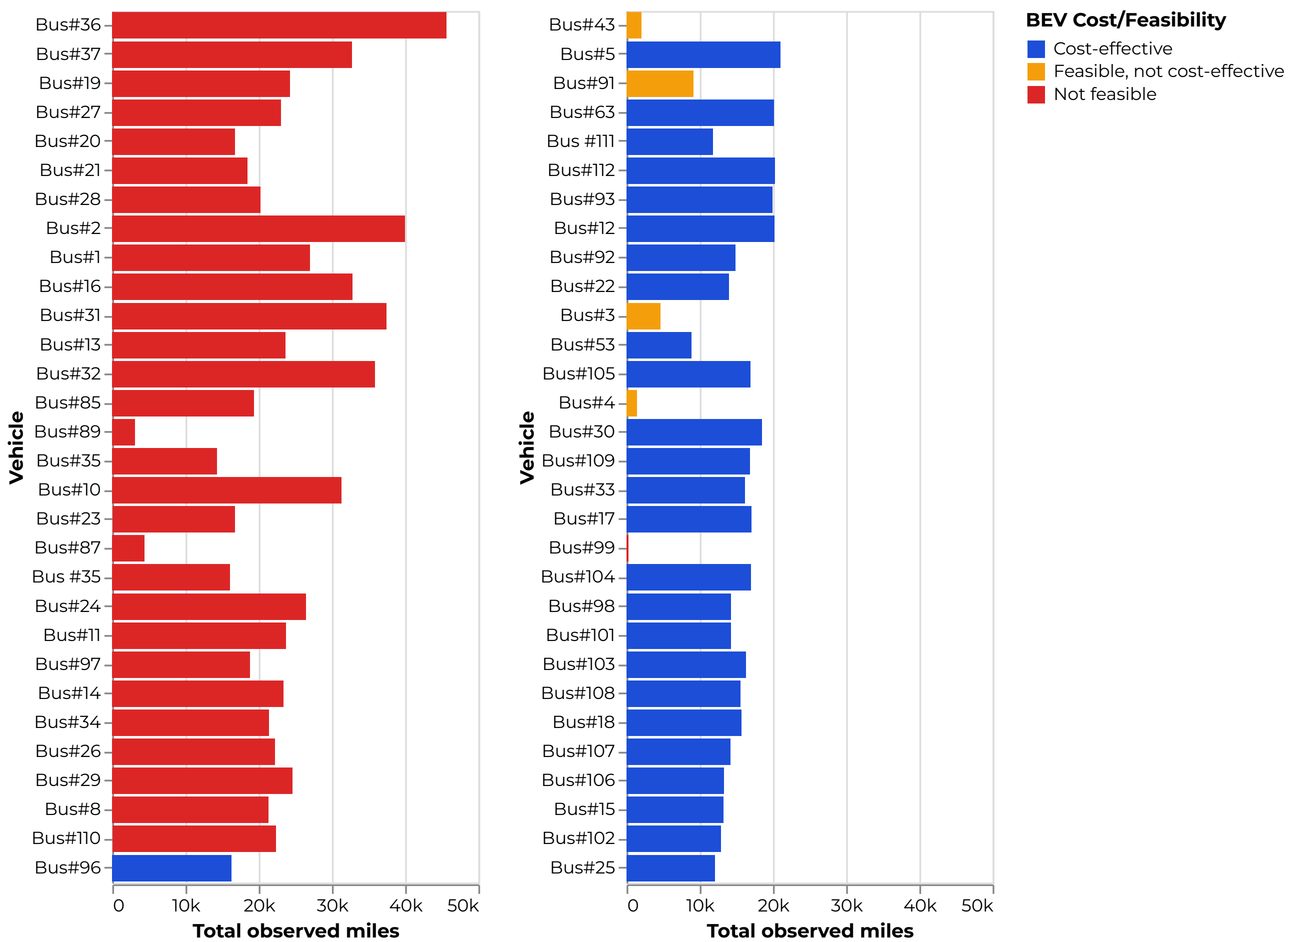

Figure 12: Total observed miles by vehicle, shown with the same ordering and cost/feasibility color classes as the required-energy diagnostic. Higher-mileage feasible buses are more likely to be cost-effective.

ments, and award amounts vary by program and funding round. Because new diesel buses would not qualify under the relevant clean-fleet programs, the modeled non-EV replacements would need to use an eligible alternative fuel.

Capex remains the dominant lever. If a district must absorb full upfront BEV purchase costs without this purchase-incentive structure, modeled opex differences alone are generally not enough to make BEV replacements lower-cost than ICEV in the near term.

4.5.3 Utilization Separates Feasible-and-Cost-Effective from Feasible-but-Not-Cost-Effective

Among buses that are technically feasible, utilization is a strong separator for economics. The feasible buses that are not cost-effective are generally those with lower lifetime miles, sporadic duty, or shorter daily use, so they accrue less fuel-cost savings to offset BEV capex.

Figure 12 makes this pattern visible directly: the "Feasible, not cost-effective" group concentrates toward lower total observed mileage, while high-utilization feasible buses are more often in the cost-effective class.

Feasibility and cost-effectiveness are both sensitive to dynamic assumptions (energy prices, vehicle purchase incentives, weather-sensitive energy burden, propulsion efficiency, and range). For that reason, the decision-support dashboard is a useful planning tool for rural districts and peer fleets: it allows transparent testing of these assumptions under a consistent modeling backbone.

5 Charging System Tradeoff Modeling

5.1 Purpose

A key limitation of the fleet-composition model is that it does not explicitly optimize charging infrastructure; charging effects are represented indirectly through energy-cost assumptions. That simplification is intentional in the dashboard stage, where fast scenario exploration is prioritized.

The purpose of charging system modeling is to resolve this limitation by moving from planning-level screening to infrastructure-level implementation analysis. In this stage, a small set of candidate fleet trajectories is selected and evaluated with detailed charging-system modeling that represents charger sizing, charger utilization, temporal charging control, and infrastructure cost implications directly. This stage also decouples charging power capacity from physical charging access: available kW determines how much power can be delivered by a charging source, while port availability determines how many buses can be simultaneously connected to a single charging source.

Methodologically, this stage is conditional on the fleet-transition scenarios rather than a separate policy exercise: fleet composition provides the BEV demand trajectory, and charging modeling tests how that trajectory can be physically and economically served over time.

5.2 Scenario Identification

To connect fleet-level economics to charging-system design, we carried three scenario definitions from the dashboard stage into the charging system optimizer. The intent is to observe how year-by-year charging operation and infrastructure decisions shift when the electrification trajectory itself changes under different economic and feasibility assumptions.

The low-, medium-, and high-electrification cases should be read as parameterized economic cases, not forecasts. Each case identifies an optimal procurement pathway under its own assumptions. The high-electrification case therefore shows more favorable overall economics partly because it assumes more favorable BEV-related conditions, including lower electricity cost, lower BEV capital cost, higher BEV purchase-incentive coverage, and improved range. Cross-scenario comparisons are useful for understanding which conditions make different electrification levels optimal, but they should not be interpreted as predictions that one pathway is more likely than another.

The scenario set spans low, medium, and high electrification conditions using different electricity/fuel prices, BEV/ICEV purchase economics, and feasibility assumptions. The fuel-price row uses a diesel planning price so the scenario table has one scalar ICEV fuel-price assumption even though the observed fleet includes both diesel and propane buses. Vehicle purchase-incentive assumptions are anchored to Utah funding programs where possible: the low-electrification case uses 63% BEV capex coverage, corresponding to a $300,000 award on a $475,000 electric school bus, and 30% ICEV capex coverage as a conservative planning assumption for eligible non-EV replacements. The medium and high cases then taper ICEV support while preserving stronger BEV support in higher-electrification conditions.

The incentive percentages in this table are scenario assumptions representing potential purchase-support coverage; they are not guaranteed program award rates.

The resulting fleet trajectories are shown in Figure 13. These trajectories are then treated as fixed annual BEV counts for the charging optimization stage.

5.3 Charging Optimization

For each fleet-year composition, the model co-optimizes charging dispatch and infrastructure investment under one objective function that internalizes multiple interacting cost streams at once: electricity tariff charges, residual liquid-fuel cost, electric vehicle supply equipment (EVSE) capital, grid-upgrade costs, battery/lifecycle costs, and other operating-cost terms. In practice, this means the solution is not just “cheapest

Table 4: Scenario assumptions passed from fleet-composition modeling into charging-system modeling.

Parameter	Low	Medium	High
Electricity price ($/kWh)	0.18	0.14	0.10
Fuel price, diesel ($/gal)	4.00	5.00	6.00
BEV capex change (%)	0%	0%	−25%
ICEV capex change (%)	0%	0%	+10%
BEV purchase-incentive coverage (%)	63%	70%	80%
ICEV purchase-incentive coverage (%)	30%	15%	0%
Electrifiable-day threshold (%)	95%	95%	95%
BEV range increase (%)	0%	20%	50%
Temperature-sensitive energy multiplier	1.0	1.0	1.0
Propulsion energy multiplier	1.0	1.0	1.0
Ambient temperature offset (°F)	0	0	0
Allowed non-depot charging stops	0	0	0
Time horizon (years)	15	15	15

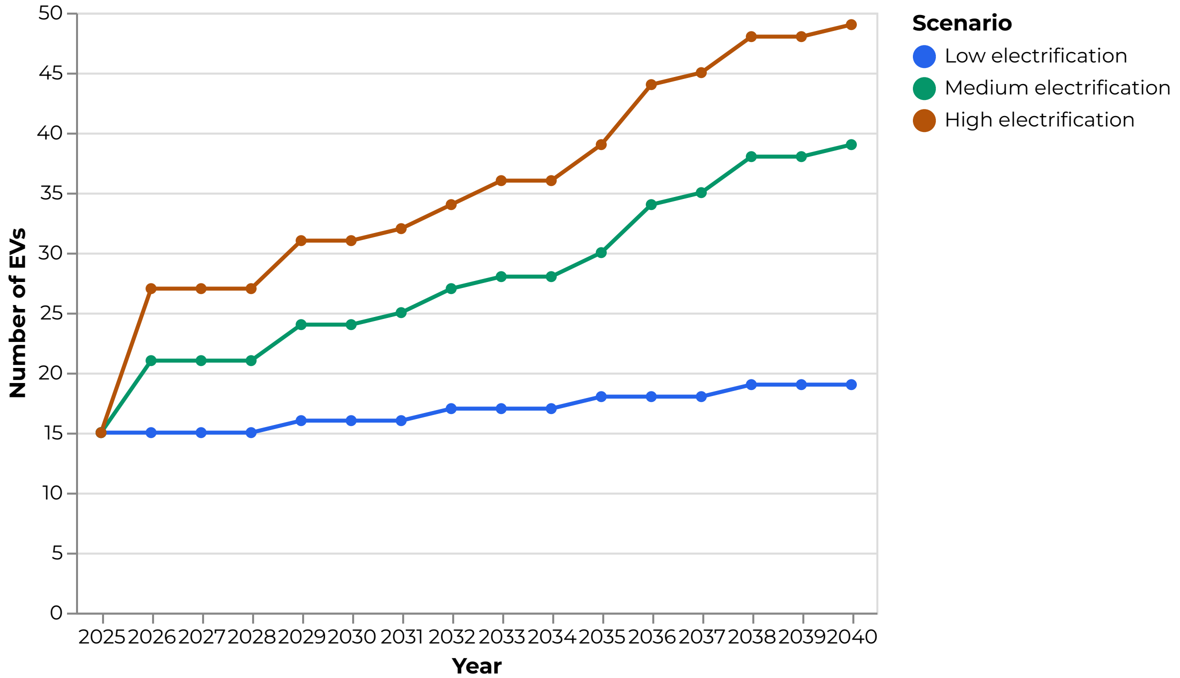

Figure 13: Scenario-specific electrification trajectories selected from the fleet-composition analysis and passed into charging-system optimization.

energy" or "fewest chargers"; it is the least-cost feasible system design after all active cost components and physical constraints are solved together.

This optimization is posed as a mixed-integer linear program (MILP) that determines how much charging to perform, where to install charging capacity, and how to satisfy vehicle energy needs at minimum total system cost for a fixed fleet-year composition.

The existing depot charging system represented in both the optimized and heuristic charging cases consists of:

- nine 16.6-kW Level 2 chargers;
- two 40-kW DC fast chargers;
- four 60-kW DC fast chargers.

Together, these chargers provide 469.4 kW of installed depot charging capacity before any modeled expansion decisions. Introducing this inventory here is important because the optimized case can use existing chargers as inherited capacity, while the heuristic case later stress-tests the same physical system under simpler rule-based operation.

For the analysis below, this inventory is represented primarily as installed charging power. The analysis does not directly model port occupancy. Instead, implementation planning should treat port availability as a deployment design requirement, with enough ports supplied for electric buses to remain connected during relevant dwell and preconditioning windows. For future electrification, this requires installation of additional ports, although not necessarily an increase to total power supply.

In practice, additional port availability would come from additional EVSE hardware, not from adding ports to the existing charging infrastructure. If a district needs more buses to remain connected while managed charging limits simultaneous power draw, the likely buildout path may be a suite of lower-power Level 2 chargers or similar port-rich infrastructure paired with controls that cap aggregate site demand.

Let decision vector $\mathbf{x}$ include charging power over time, infrastructure-sizing variables, and discrete infrastructure-selection variables. The optimization objective is:

$$\begin{aligned}\min_{\mathbf{x}} C_{\text{tot}} = C_{\text{elec}} + C_{\text{fuel}} + C_{\text{evse}} + C_{\text{grid}} \\ + C_{\text{batt}} + C_{\text{repl}} + C_{\text{opex}}\end{aligned}$$

where each term is a net-present-cost component (electricity tariff cost, liquid fuel cost for non-electrified operation, charging equipment, grid-upgrade costs, battery costs, replacement costs, and other operating costs).

Core feasibility constraints are linear and include the following relationships.

Charging energy is assigned to electrified vehicles according to their required energy:

$$\sum_{l,t} \eta\, \Delta T\, p_{vlt} = y_v\, E_v^{\text{req}} \quad \forall v$$

Vehicle state of energy evolves from charging and driving energy use:

$$s_{v,t+1} = s_{vt} + \eta\, \Delta T \sum_{l} p_{vlt} - e_{vt} \quad \forall v, t$$

State of energy must remain within allowable battery bounds:

$$s_v^{min} \le s_{vt} \le s_v^{max} \quad \forall v, t$$

Total charging power at each charging location must stay within available power capacity:

$$\sum_{v} p_{vlt} \le K_l \quad \forall l, t$$

Grid load must remain within the modeled grid envelope, plus any allowed upgrade or excess-capacity variable:

$$P_{gt}^{\text{net}} \le \bar{P}_{gt} + u_{gt}, \quad u_{gt} \ge 0 \quad \forall g, t$$

and electrification constraints fixing each vehicle to the scenario-specified fuel state:

$$y_v = \begin{cases} 1, v \in \mathcal{V}_{\text{BEV},y} \\ 0, v \notin \mathcal{V}_{\text{BEV},y} \end{cases}$$

where $\mathcal{V}_{\text{BEV},y}$ is the BEV set for modeled year y.

Within the time horizon of an electrification trajectory, existing infrastructure is carried forward so previously built capacity is available in subsequent years:

$$K_{l,y} \ge K_{l,y-1}$$

where $K_{l,y-1}$ is the optimized charging capacity already available at location l from the prior modeled year. Grid-envelope upgrades follow the same carry-forward logic: once added, the upgraded grid capacity is treated as available in subsequent modeled years rather than re-purchased from scratch.

Term definitions used above:

- v, l, g, t: vehicle, charging location, grid asset, and timestep indices.
- p_{vlt}: charging power assigned to vehicle v at location l and time t (kW).
- η: charging efficiency.
- ΔT: timestep duration (hours).
- E_v^{req}: total traction-energy requirement for vehicle v over the modeled period (kWh).
- y_v: binary electrification indicator for vehicle v (1 = BEV, 0 = non-BEV).
- s_{vt}: vehicle state of energy at time t (kWh).
- e_{vt}: driving-energy draw assigned to vehicle v at time t (kWh).
- s_v^{min}, s_v^{max}: minimum and maximum allowable state of energy (kWh).
- K_l: available charging power capacity at location l (kW or equivalent charger capacity).
- P_{gt}^{net}: modeled net grid demand on grid asset g at time t (kW).
- $\bar{P}_{gt}$: allowed grid-envelope capacity for asset g at time t (kW).
- u_{gt}: nonnegative excess above grid envelope (kW), priced in cost terms when enabled.
- $K_{l,y-1}$: charging capacity carried forward from the previous modeled year.

This is a power-capacity constraint. It limits aggregate kW delivered at a charging location; it does not directly constrain or optimize the number of buses physically connected at once. The ports-versus-power interpretation therefore depends on a separate implementation assumption: sufficient physical ports can be supplied for buses that need to remain connected, while managed charging controls how much power those connected buses draw.

Reported optimization outputs for each solved year include:

- optimized charging schedules by vehicle, location, and timestep;
- installed charging power/capacity by site, including whether capacity is inherited or newly added;
- energy throughput and utilization summaries by site and charging paradigm, with port availability treated as a separate implementation requirement;
- electricity-bill components (energy, demand, and time-varying demand where applicable);
- cost decomposition consistent with C_{tot} and solver diagnostics (objective value, bound, optimality gap, runtime, and model size).

5.4 Observations

5.4.1 Depot-Only Charging Covers the Modeled BEV Demand

Depot-only charging is a deliberate modeling constraint in this case study: we intentionally require charging to occur at district-controlled depot locations, rather than assuming dependable public or third-party charging access. That framing matches school-district decision authority and avoids embedding uncertain external infrastructure expansion in the core planning result.

For most buses in this dataset, this outcome is directionally expected because vehicles spend substantial portions of the school-day cycle parked at depot between service windows. The optimizer therefore finds technically feasible charging plans without relying on non-depot charging for the buses included in each trajectory.

For the subset of long-distance duty cycles that remain difficult to electrify, non-depot charging would likely be the enabling mechanism. However, that pathway is outside a school district's direct control and depends on uncertain public-charging development timelines, access rules, and reliability.

The optimized buildout result shows no additional grid power capacity beyond the existing depot inventory in any of the low-, medium-, or high-electrification pathways, though for operations and energy management reasons adding ports on existing grid capacity will likely be beneficial. In all three pathways, the model serves the selected BEV fleet by coordinating use of the district's existing Level 2, 40-kW DC fast, and 60-kW DC fast charging capacity. This should be read as a modeled power/capacity result: it does not mean physical connection points are unim-

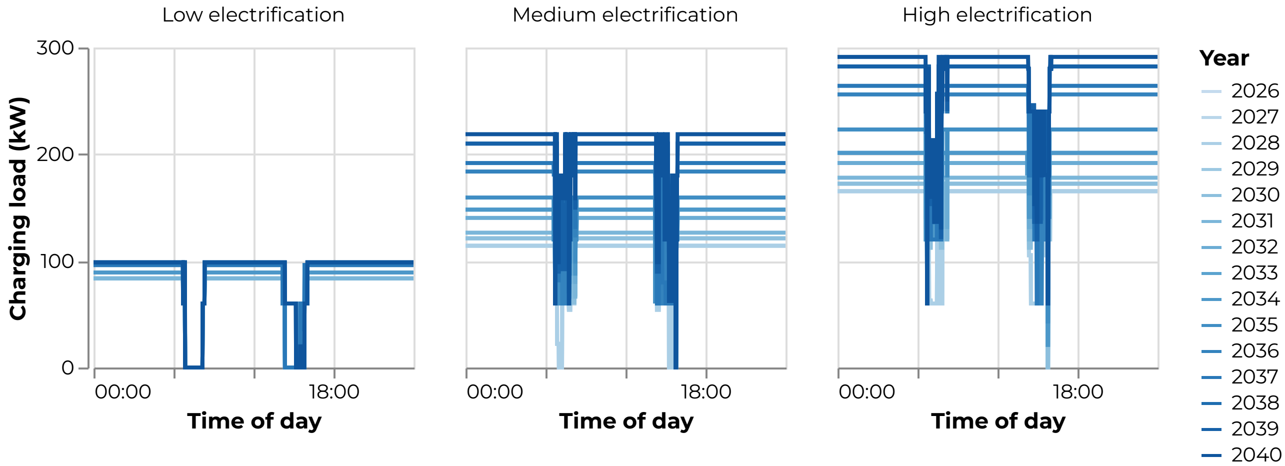

Figure 14: Average daily charging load profile by scenario and year, shown on a normalized daily clock axis.

portant. Instead, it indicates that the need for more ports should be evaluated separately from the need for more simultaneous charging power.

5.4.2 Managed Charging Helps Share Power Capacity

The optimizer assumes idealized operational coordination: charging power can be shared without administrative friction, and no additional schedule buffers are imposed beyond the modeled constraints. Under that assumption, these solutions should be interpreted as a lower bound on achievable charging-system cost for each electrification trajectory.

The value of the optimized result is therefore not the exact hour-by-hour schedule. It identifies the benefit of coordinating charging power, avoiding unnecessary peaks, and using dwell windows effectively. A field implementation should deliberately add operating margin around that benchmark so a delayed return, missed charge, unavailable charger, or communication failure does not cascade into a route-readiness problem.

The schedule structure in Figure 14 illustrates why shared and coordinated infrastructure can outperform heuristics such as dedicating one charging power source to every bus and charging immediately upon arrival. Charging is shifted in time to exploit dwell windows and tariff structure, and interconnect utilization increases materially when charging demand rotates between buses using shared assets.

In this simulation, the model allows extensive sharing among EVSE and buses and uses highly coordinated tactics to find the least-cost feasible operation, so it can concentrate charging on flexible higher-power assets in ways that may be undesirable or difficult to reproduce as a field operating strategy. Charging profiles that are similarly responsive to electricity-rate incentives, but allow buffer time and redundancy for resilience, should be a practical target for managing operations.

This effect remains important even though the optimized run does not select incremental charger buildout. The model is able to serve higher BEV shares by using the existing shared depot chargers more intensively, rather than by adding a dedicated charger for each bus.

The planning implication is that power capacity and port availability are different design choices. Managed charging can share available kW across connected buses, while adequate port access lets buses remain connected during dwell and preconditioning windows without repeated vehicle moves. This is why the optimized result should be paired with a field plan for port access, charger uptime, and recovery margin rather than read as a recommendation to minimize connection opportunities.

5.4.3 Coordinated Charging Lowers Delivered Charging Cost

In the fleet-replacement analysis, electricity price in $/kWh is treated as a scalar planning input. In the charging-system model, electricity cost is calculated from the modeled utility tariff instead: charging load is assessed both energy costs and peak demand costs according to the source district's current rate schedule (Rocky Mountain Power Schedule 6). The scalar dashboard price and the optimized tariff bill therefore answer related but different questions. The dashboard scalar price supports fast fleet-economics sensitivity testing, while the charging model estimates the bill created by a specific managed charging schedule under a utility rate structure.

The charging-system model lets us deconstruct the scalar planning assumption into explicit annual cash flows, discounted net-present-cost (NPC) components, and tariff electricity-bill components tied to the realized electrification trajectory. The next two figures broaden the view from electricity billing alone to total system cost: they include vehicle purchases, charging-infrastructure costs, diesel fuel costs for remaining non-BEV operation, and electricity costs for BEV operation. Because the low-, medium-, and high-electrification cases use different vehicle-cost, incentive, fuel-price, and electricity-price assumptions, their total costs should be interpreted as internally optimized pathways under different parameter conditions rather than as a direct forecast ranking.

The progression from operations to economics is:

1. Build annual undiscounted cash flows by cost component (infrastructure installs/replacements and annual opex), shown in Figure 15.
2. Discount those annual values to the analysis base year and sum by component to get NPC, shown in Figure 16.
3. Divide discounted tariff electricity-bill components by discounted delivered charging energy to get electricity-bill cost per delivered kWh, shown in Figure 17.

This decomposition shows that the high-electrification scenario achieves approximately $0.063/kWh optimized tariff electricity cost under the modeled assumptions, which is significantly below the baseline $0.18/kWh planning value used in the dashboard analysis. That result should be interpreted within the favorable high-electrification parameter set, not as a prediction that the district will face that tariff cost or reach that fleet composition. It indicates the electricity-bill cost floor implied by rate-responsive allocation of charging power in an optimized depot-charging design.

The demand-charge component in Figure 17 should be read as cost per delivered kWh, not as total demand-charge dollars. Demand-related costs typically increase in absolute terms as electrification grows and charging peaks become larger. However, higher-electrification cases also deliver much more energy, so higher total demand charges can still appear as a similar or lower demand-charge cost per delivered kWh when they are spread across a larger energy denominator.

5.5 Heuristic Infrastructure Usage Stress Test

The optimized charging case above represents a highly coordinated design: charging schedules, power allocation, charger use, and infrastructure decisions are selected together to minimize total system cost while adhering to the modeled operating constraints. That result is useful as a cost-efficient planning benchmark, but it is not the only relevant implementation question for a district fleet.

The heuristic charging case asks a more operational question: if buses follow a simple rule-based charging policy, how hard can the existing depot charging system be worked, and what tradeoffs appear? In this case, buses connect when they return to a charging location, charge when compatible EVSE power is available, and wait in a queue when the port they are connected to is not receiving power because the associated EVSE is serving another bus or because power is curtailed to limit aggregate demand. This provides an operational stress test of existing infrastructure under less centralized coordination.

The two cases should be interpreted together. The optimized case estimates what is possible under coordinated power allocation and system design. The heuristic case estimates what may happen under a simpler operating policy where charger use and queuing pressure become visible. The comparison helps distinguish charging-power needs from management needs such as

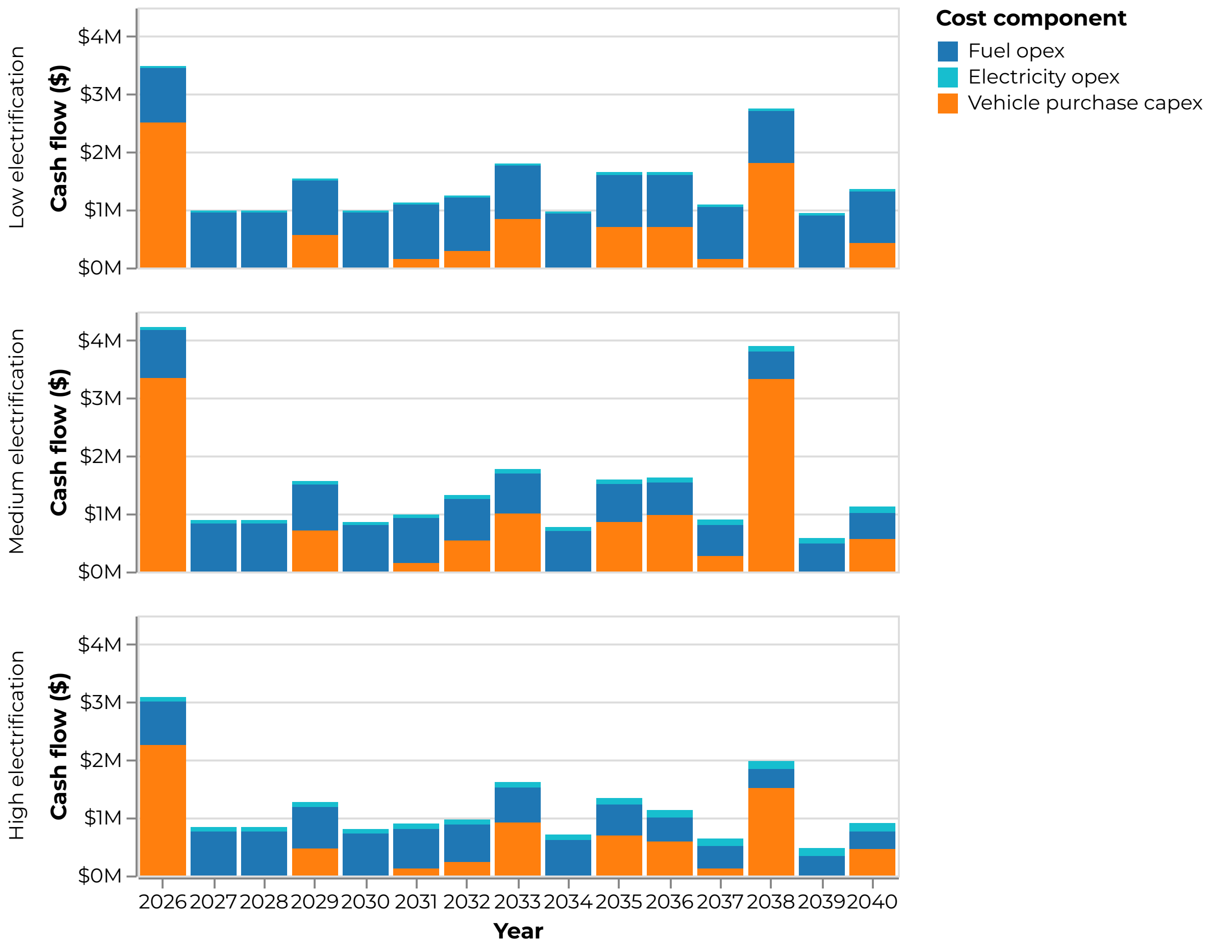

Figure 15: Undiscounted annual cash flows by component for each charging-system scenario.

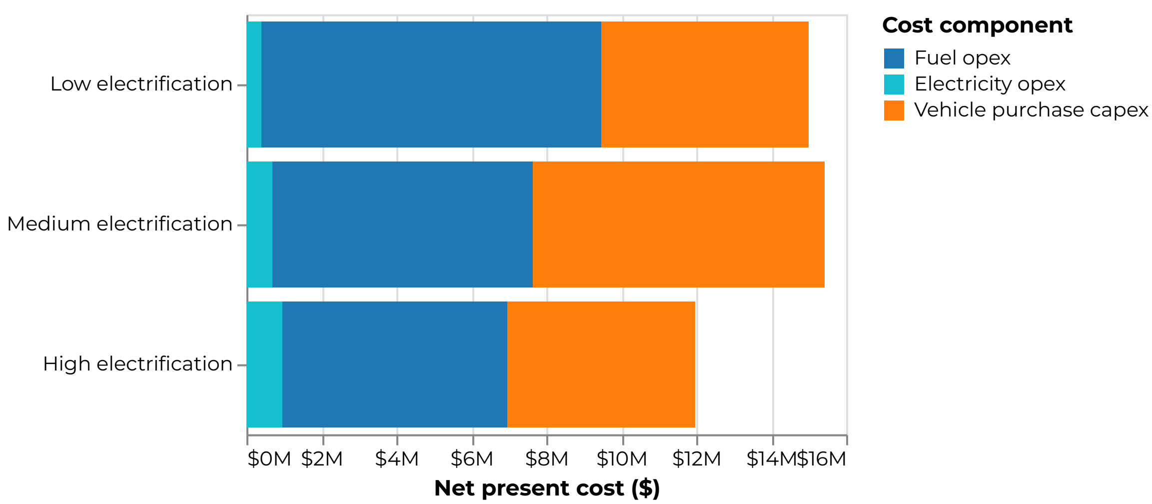

Figure 16: Net present cost decomposition by component and scenario.

port availability assumptions, queuing, peak-demand control, vehicle prioritization, and charger uptime.

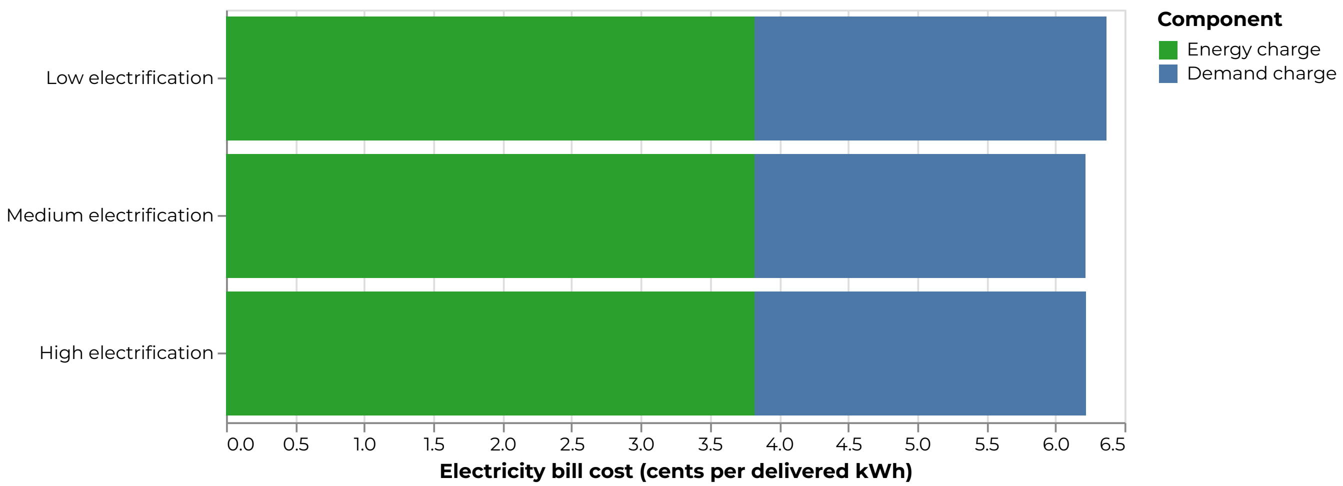

Figure 17: Optimized-case electricity bill cost by scenario and cost component, expressed in cents per delivered kWh.

5.5.1 Method

The heuristic case uses the same low-, medium-, and high-electrification fleet trajectories evaluated in the optimized charging case. It also uses the same underlying trip-energy demand and depot-only charging frame, but simulates the full available trip record rather than optimizing a short representative window.

The heuristic simulation uses the existing depot charging inventory introduced in Section 5: nine 16.6-kW Level 2 chargers, two 40-kW DC fast chargers, and four 60-kW DC fast chargers. Vehicles are assumed to begin with a full usable battery, usable battery capacity is set to 80% of nominal capacity, charging efficiency is set to 90%, and vehicle state of charge carries across days. A five-minute EVSE turnover time is imposed between charging sessions. Electricity-cost calculations use the same utility rate schedule as the optimized charging case.

This case prioritizes charging the lowest-state-of-charge buses first. In a production charging-control system, priority should also account for the next duty cycle: departure time, expected trip energy, reserve target, preconditioning need, and whether a missed charge would create a recovery problem for later service.

The heuristic case begins with existing charging infrastructure and allows additional EVSE units to be added only if the simulation identifies a correctable infeasibility and a least-cost feasible charger addition is selected. In the modeled district run, no EVSE additions are selected by this expansion process; the heuristic case therefore evaluates how the existing power supply performs as BEV operation increases (assuming, but not directly modeling, sufficient port availability).

Queuing is central to the heuristic case, but it should not be read as a physical line of buses waiting for an open port. In this model interpretation, a queued vehicle is already connected to a compatible port, but the EVSE associated with that port is supplying power to another connected bus (one-bus-at-a-time charging for each EVSE is a modeling simplification; allocating power between connected buses in practice would provide additional charging-management flexibility). When a vehicle arrives below its target state of charge and compatible EVSE power is available, it begins charging. When compatible EVSE power is assigned elsewhere, the connected vehicle waits in a queue. If a vehicle departs while charging, the session ends at departure. If a vehicle departs while queued, it disconnects and begins its next trip with its current state of charge.

This structure is intentionally simple. It does not search for a globally optimal future charging plan, avoid demand-charge peaks in advance, or reassign charging sessions to optimize systemwide cost. Its purpose is to represent a plausible charge-management rule and reveal the operational consequences of relying on existing shared EVSE. In this stress test, queuing is a proxy

Table 5: Optimized and heuristic charging-method comparison.

Modeling dimension	Optimized charging case	Heuristic charging case
Main question	What charging-system design minimizes total cost under coordinated control?	How does the existing charging system perform under simple rule-based operation?
Charging dispatch	Globally coordinated power allocation across vehicles, EVSE, and timesteps	Connect on arrival; charge when compatible EVSE power is assigned; queue while connected but not receiving power
Infrastructure treatment	Infrastructure and charging operation are considered jointly	Starts from existing infrastructure; adds EVSE units only if the expansion logic requires it
Operational coordination	Idealized coordination and managed power sharing	Local rule-based EVSE power assignment and queue priority
Time scope	Representative optimized windows for selected fleet years	Full available trip-energy input period for each modeled fleet year
Best interpreted as	Lower-bound coordinated design cost	Operational stress test

for shared EVSE power pressure among already connected buses, not a direct model of physical port scarcity and not simply a lack of total energy or aggregate kW over the full day.

5.5.2 Observations

5.5.2.1 Existing Depot Power Capacity Can Serve More Energy Under High Utilization

The heuristic case does not add grid capacity beyond the existing depot system in the modeled run. This is an important result, but it should be interpreted carefully. It does not mean the existing system has large spare capacity or that physical port access is unconstrained. Rather, it shows that the existing system is large enough to absorb substantial additional BEV charging demand when port availability is sufficient and vehicles are automatically managed to share available EVSE power, wait while power is assigned to other connected buses, and charge whenever charging windows are available.

For planning, the result shifts attention from whether the district can use the existing charging system to what operating discipline would be required at higher utilization. The same grid infrastructure can be technically capable of serving the energy demand while still requiring a separate port-availability plan, staff or software coordination, readiness buffer, and careful management to avoid missed charging windows, prevent unnecessary peaks, and ensure the right buses receive priority.

The readiness buffer is especially important in winter. If a bus cannot remain connected before departure, the fleet may lose the ability to precondition automatically, and the same route can require more battery energy once cabin heating, battery conditioning, and cold-weather performance effects are active. For that reason, connected parking and managed power allocation should be treated as reliability tools, not only cost-control tools.

5.5.2.2 Queuing and Peak Load Are the Main Tradeoffs

The heuristic case relies on EVSE power sharing and queuing. As electrification increases, buses ideally have access to dedicated charging ports, but share existing grid capacity through a digitally managed shared allocation or queuing program. This flexibility allows automated man-

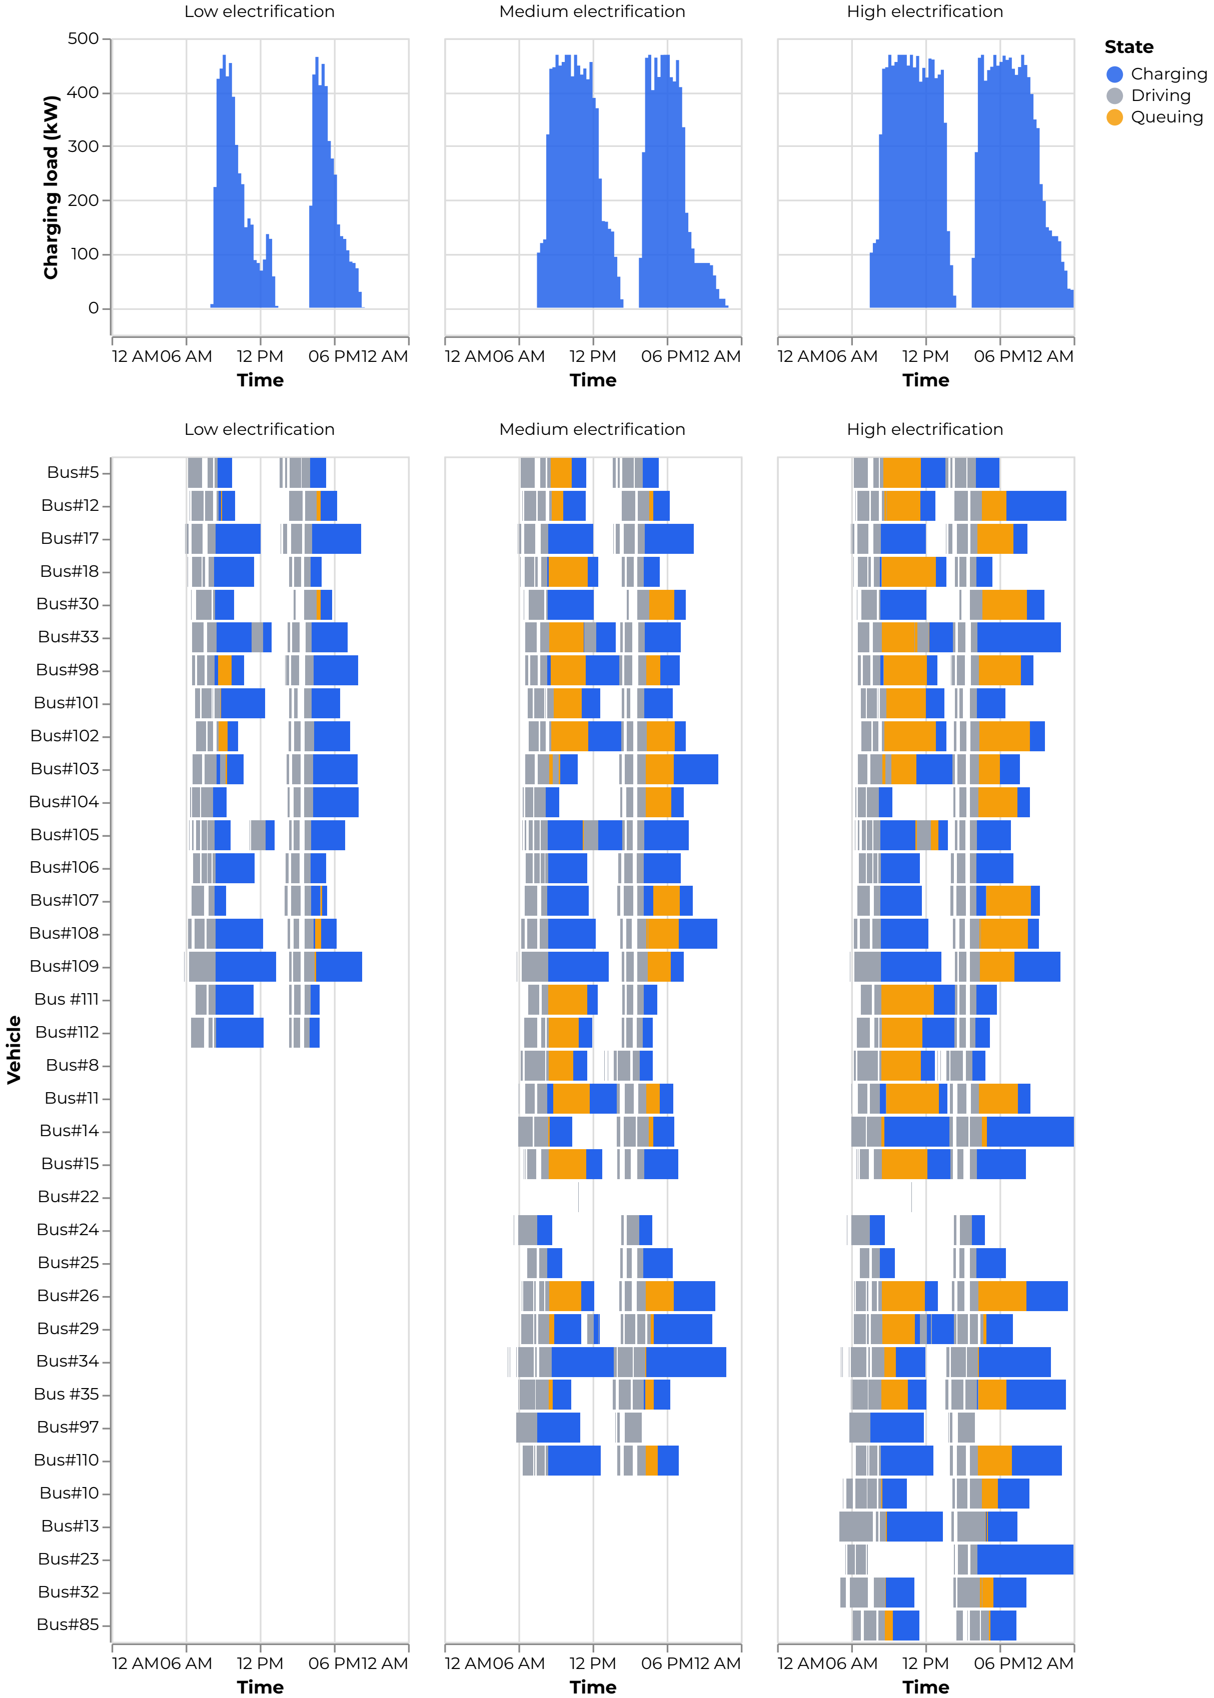

Figure 18: Representative final-year heuristic-case operation, showing total charging load above bus-level driving, charging, and queuing intervals.

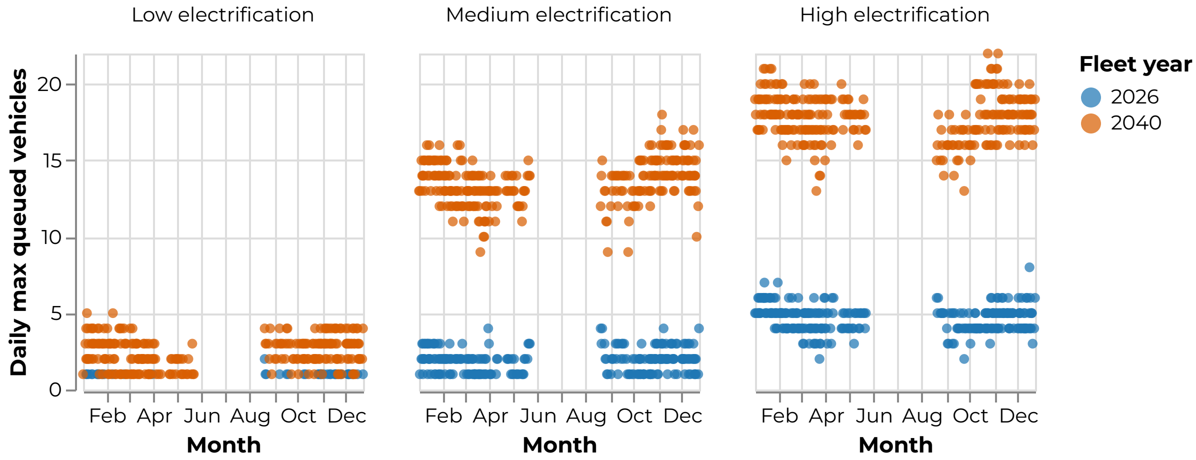

Figure 19: Heuristic-case queue pressure across the operating domain, showing daily maximum queued vehicles by scenario for the initial and final modeled fleet years.

agement of peak demand and other utility-rate incentive responses.

Figure 18 shows this interaction for a representative final-year day. The upper panel shows total charging load over time, while the lower panel shows vehicle activity intervals. Charging, driving, and queuing are shown as distinct activity states. The figure therefore emphasizes the operational pattern that matters for this discussion: connected buses can rotate through available power over the day, but higher BEV demand also increases the importance of port access, automated prioritization, and recovery margin.

Shared EVSE operation increases utilization, but without coordinated management it also creates peak-load exposure and resiliency vulnerabilities. In the heuristic output, final-year peak charging load reaches approximately the installed depot capacity in the medium- and high-electrification scenarios. That indicates very high use of available depot power during peak periods. It also means peak-demand exposure remains a meaningful cost and utility-planning concern.

In this context, queuing means the vehicle needed energy, was connected at a charging location, and waited because compatible EVSE power was assigned to another connected bus or temporarily unavailable due to turnover. Some queuing is a natural consequence of shared power infrastructure. However, persistent or poorly timed queuing would create real operational burdens: charging priority may need to be monitored, staff or software may need to ensure that critical vehicles do not miss charging windows, and the depot needs enough ports for queued buses to remain connected while waiting for power.

The result illustrates why port availability matters separately from power capacity: a system can have enough aggregate kW to deliver energy while still creating operational pressure if too few buses can remain connected during dwell and preconditioning windows. Because queued buses are assumed to remain connected while waiting for EVSE power, adequate port availability is part of making this shared-power strategy workable. The key design goal is to preserve the benefits of managed power allocation while reducing manual movement, missed-charge exposure, and queue risk.

Figure 19 summarizes daily maximum queue depth across the modeled period. This diagnostic is useful because the representative-day figure can show how queuing works operationally, while the domain-level view shows whether queuing is isolated or recurring.

The heuristic case also assumes that the charging system remains reliably operated: EVSE are available when modeled, no unscheduled outages occur, and automated charge management continues well beyond school hours when needed. That assumption is optimistic for real deployments. If EVSE downtime, insufficient port availability, communication failures, late vehicle

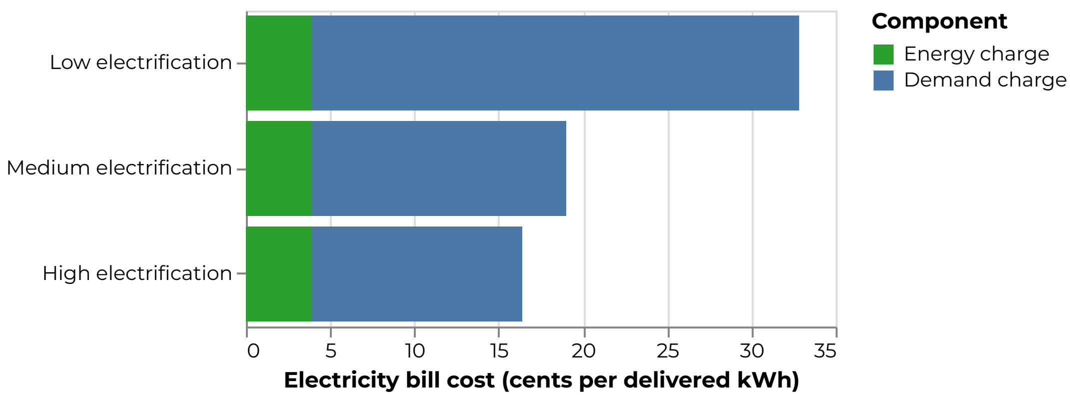

Figure 20: Heuristic-case electricity bill cost by scenario and cost component, expressed in cents per delivered kWh.

returns, or missed connections are common, the same nominal charging inventory could perform materially worse.

A functional deployment should therefore include redundancy and recovery time as design requirements. That may mean keeping spare port and EVSE availability, preserving some unused power capacity for recovery, setting minimum state-of-charge buffers above the modeled feasibility floor, flagging missed or delayed charging sessions quickly, and using operating rules that prioritize buses whose next duty cycle has the least recovery room.

This simulated future-load result should be read alongside the observed-load analysis in Section 6.1. The heuristic simulation represents future electrification pressure under rule-based charging and produces higher peak loads than the district's currently observed charging behavior. That difference is not a contradiction; it reflects the gap between today's partial EV fleet operation and a modeled future in which many more buses draw on the same shared depot charging system.

5.5.2.3 Higher Utilization Can Lower Electricity Bill Cost per Delivered Energy

The economics of the heuristic case illustrate how unsophisticated charge management can increase electricity cost burdens. However, as more BEV energy is delivered through the same existing grid capacity, fixed and peak-related costs are spread over more delivered energy.

For this reason, the heuristic case shows lower electricity bill cost per delivered energy in the higher-electrification scenarios, even though those scenarios create more operational intensity. In round terms, the final-year high-electrification heuristic case delivers more than 1 GWh of energy after charging efficiency losses annually through the existing depot charging inventory, and its electricity bill cost per delivered kWh is lower than in the lower-electrification case. This is an economic utilization result within the high-electrification parameter case, not a forecast or complete operating recommendation; it should be read alongside queue depth, port availability, staff effort, and route-readiness buffer.

Figure 20 shows the electricity bill cost components for the heuristic case. The values should be interpreted as cost per delivered kWh. This denominator matches the planning question of how much usable energy is delivered to the fleet, while still preserving the effect of charging losses through the underlying electricity bill calculation.

Figure 21 compares the levelized tariff electricity costs between the optimized and heuristic cases. These two cases should be interpreted as planning reference points rather than as equally likely operating plans. The optimized case represents a highly coordinated lower-bound benchmark: charging is scheduled to satisfy vehicle energy needs while responding directly to rate-schedule incentives, including demand-charge and time-of-use exposure. That level of coordination is likely possible with a well-designed charge management system, but it depends on strong software integration, operational discipline, recovery

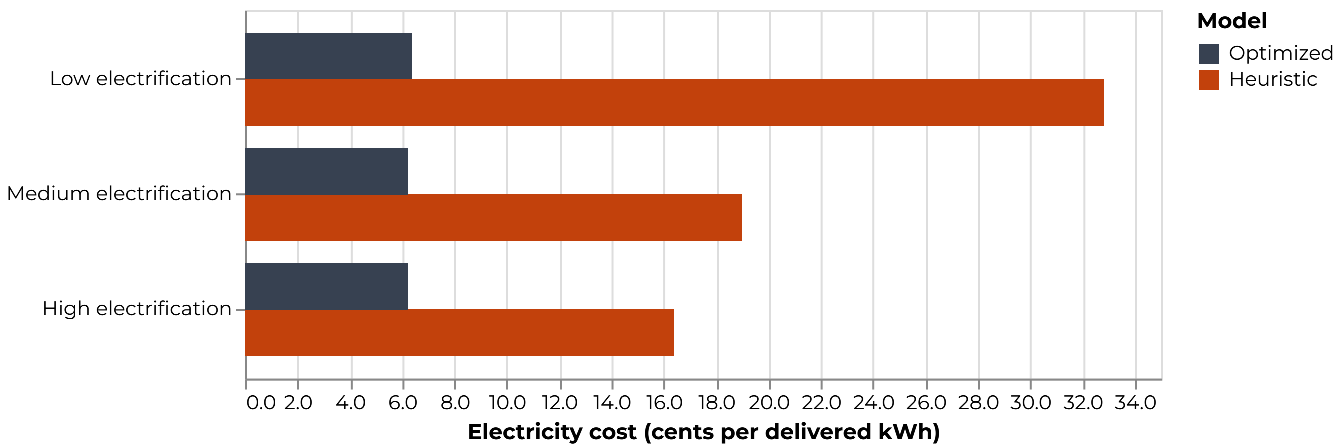

Figure 21: Electricity cost for optimized and heuristic charging cases, expressed in cents per delivered kWh.

buffer, and EVSE/port availability; in addition, resiliency requirements such as recovery buffer and outage management necessitate some tradeoff in terms of higher capital and electricity costs. The heuristic case represents a simpler charge-on-arrival policy that is easy to understand but suboptimal from a utility-bill perspective because it concentrates load on the existing depot EVSE and creates larger demand-charge exposure. It is also optimistic operationally because it assumes reliable EVSE availability and successful execution of the queuing/control rule.

The planning takeaway is that energy management can materially improve operating cost by moving real-world charging behavior away from the heuristic reference point and toward the rate-responsive optimized benchmark. A district does not need to reproduce the exact optimized schedule for energy management to be valuable. Even partial coordination–such as delaying non-urgent charging, avoiding tariff peak windows where possible, staggering charger starts, and prioritizing buses with near-term service needs–can reduce electricity bill cost by making charging respond to the incentives embedded in the rate schedule. The operating goal should be a robust managed-charging strategy that preserves dispatch readiness first, then uses the remaining flexibility to reduce peaks and energy cost. Utilization still improves within the heuristic case as electrification rises, which is why the heuristic cents/kWh falls from the low to high pathway, but it remains above the coordinated optimized benchmark.

The magnitude of demand-related costs depends strongly on the rate schedule used for this calculation. That rate sensitivity is the bridge to the final section, which turns from modeled future fleet trajectories to the source district's observed charging load.

The optimized-charging and heuristic results above ask how the depot charging system could perform as electrification grows. The following load-profile analysis asks a more immediate implementation question: what does the district's existing EV charging behavior imply for utility-rate selection, billing risk, and smart energy management? It uses the current Fleetbox charging records to compare the current Schedule 6 billing estimate with Schedule 6A as a near-term rate opportunity, while also using Schedule 8 as a future fleet-expansion and rate-design comparison.

6 Charging System Solution Design Considerations

6.1 Present Charging Patterns, Rates, and Smart Energy Management

6.1.1 Executive Summary

The source district's electric buses create a highly structured charging load: most charging occurs after buses return from morning pickup routes and again after afternoon drop-off routes. As described in Section 6.1.3.1, that one-year load profile contains about 160 MWh of charging energy. Because the charging is concentrated in a few repeatable windows, the tariff applied to that load matters as much as the total energy consumed. Under the current tax-exempt Schedule 6 billing estimate, the observed charging profile costs about $55,200 per year. Applying the same charging profile to Schedule 6A reduces the annual cost to about $36,700, a savings of about $18,500, or 34%, with no operational change.

The analysis also shows meaningful opportunities for smart energy management, as described in Section 6.1.3.3. In this report, smart energy management means the buses still get the charging energy they need on their normal operating cadence, while grid costs and demand windows help determine the exact charging hours within the windows the buses already use. Practical strategies include shifting energy away from high-cost time windows and flattening monthly demand peaks when doing so does not compromise bus readiness. A modeled Schedule 6A smart energy management case reduces the annual cost to about $26,000, bringing total savings relative to the current Schedule 6 baseline to about $29,200, or 53%. A modeled Schedule 8 smart energy management case is lower at about $14,800 per year, a savings of about $40,400, or 73%. However, the district's observed load does not naturally meet the Schedule 8 applicability threshold. The Schedule 8 result should therefore be read as a future fleet-expansion and rate-design signal rather than as an immediately available tariff option.

These savings estimates are large enough to support action, but the source data has limitations and is not a revenue-grade meter record. As discussed in Section 6.1.3.1, the 15-minute power data and hourly energy reports broadly align with expected charging patterns. However, the monthly peak demand from the hourly energy reports is higher than the peak in the 15-minute data for each overlapping month evaluated. Undocumented sampling, rolling, rounding, filtering, or missing intervals in the 15-minute endpoint could account for some or all of this discrepancy. The load profile is still useful for identifying charging patterns, rate opportunities, and smart energy management strategies, but implementation decisions should first cross-check the Fleetbox-derived monthly peaks against the monthly billing peak demands and, where available, Rocky Mountain Power 15-minute interval or load-profile data.

6.1.2 Methodology

6.1.2.1 Fleetbox Data Extraction

Fleetbox is the district's charging management and data platform. The source district provided ASPIRE access to Fleetbox for this analysis, but Fleetbox did not provide a complete export feature for the hourly and sub-hourly charging records needed for this analysis, nor for charging-session-defined data. To obtain the data, the web application's normal browser requests were inspected so the same authenticated API requests used by the site could be documented and repeated. Those requests were then executed from a browser user-automation framework, preserving the normal authenticated website context.

Two Fleetbox data products were used. The first was the hourly energy report, which reports imported and exported energy by EVSE, or charger. The hourly charger-level report was extracted through both the report-export workflow and the API request path. The API export was used as the primary source because it preserved timestamp and time-zone information, including daylight-saving transitions. The report-export path treated all local days as 24-hour days and did not reliably represent 23-hour or 25-hour daylight-saving transition days. The API records carried charger identifiers, and the charger-level report-export metadata was needed to map those identifiers back to human-readable charger names and serial numbers.

The second data product was the site-level real-time power endpoint, requested at a 900-second interval (15 minutes). The endpoint returned a single total active power time series rather than charger-level records. The Fleetbox website exposes this endpoint interactively at five-minute resolution and without the same long-range history available from the hourly energy reports. Repeating the authenticated API request at the 15-minute interval made it possible to evaluate the shape of total charging demand at billing-demand resolution for the available overlap period and to compare monthly peak demand against the hourly energy reports. All timestamps were converted to Mountain Time so that charging patterns align with local school operations and Rocky Mountain Power billing windows.

6.1.2.2 Rate Schedule Billing

The billing comparison uses the hourly energy profile from May 1, 2025 through April 30, 2026. Charger-level records are aggregated to one site-level row per hour. Because each row represents one hour, hourly energy in kWh is numerically equal to average power in kW for that hour. The same site load is then billed under Rocky Mountain Power Schedule 6, Schedule 6A, and Schedule 8.

All bill values are modeled tariff simulations applied to past charging use. They should be read as estimates of what future bills could look like if the same EV charging load occurred under the modeled tariffs and billing assumptions. They are not a reconstruction of actual historical Rocky Mountain Power invoices. In particular, the observed Schedule 6 estimate may not match actual bills for the May 2025 through April 2026 period because tariffs, riders, or other bill components may have changed since those bills were issued, and because the Fleetbox source data is not the utility's revenue meter.

The modeled site load is also limited to the EV charging load captured in Fleetbox. If other building, depot, maintenance, or temporary loads are connected to the same utility meter now or in the future, their energy use, peak demand, and timing would need to be included before using these estimates for billing or infrastructure decisions.

The rate model includes fixed customer charges, energy charges, demand charges, facility charges, applicable riders, and the paperless billing credit. State and local sales taxes are set to zero because the source district is a government entity. For tariffs with on-peak periods, the model separates on-peak and off-peak energy and demand using the schedule's month-specific time windows, while treating weekends and observed holidays as off-peak. Schedule 6 is treated as the current baseline. Schedule 6A is treated as an available alternative rate that the district could choose to utilize while its load remains below the Schedule 8 applicability threshold. Schedule 8 is modeled as a potential future-value case for a larger load profile that is subject to, or can qualify for, Schedule 8[1].

6.1.2.3 Smart Energy Management Model

Smart energy management was modeled as an hourly linear program that reshapes charging energy without changing the daily, overnight, or weekend service requirements. This is a screening-level modeling estimate, not a dispatch-ready operating plan. The model preserves energy in broad weekday daytime, weeknight, and weekend buckets, but it does not model bus-by-

[1] Rocky Mountain Power Schedule 8 applies after load has registered at least 1 MW more than once in the preceding 18 months, and remains applicable until a subsequent 18-month period passes without any 1 MW-or-higher demand reading. The source district's current load profiles do not approach 1 MW; the site's peak demand under scheduled bus operations is much lower, and even the potential charging peak from the chargers currently on site is less than 1 MW. The district reported that Rocky Mountain Power identified a 1,200 kVA interconnect, so utility-side upgrades may not be needed for a future Schedule 8-scale load, but the charging equipment and operating load would still need to grow substantially. The published schedule does not specify whether the initial "more than once" requirement must occur in separate months or can be satisfied by multiple billing intervals. Any intentional operating change intended to affect Schedule 8 applicability should be evaluated only in close coordination with Rocky Mountain Power and with review of the district's utility contract, service agreements, and applicable legal or procurement requirements.

bus state of charge, charger assignment, connector occupancy, charge acceptance curves, route-specific readiness, communications or control failures, or other field constraints that a production control system would need to handle. Let x_t be managed charging energy in hour t, and let p_m be the modeled monthly peak for month m. The objective minimizes the sum of monthly peaks, with a small timing penalty that chooses earlier charging when multiple peak-minimizing solutions are equivalent:

$$\min_{x_t, p_m} \sum_m p_m + \epsilon \frac{\sum_t s_t x_t}{\sum_t e_t}, \tag{1}$$

where e_t is observed hourly charging energy, s_t is a sawtooth timing penalty within each flexible charging bucket, and ϵ is small relative to the monthly peak objective.

Energy is preserved within operationally meaningful buckets rather than across the full year. Weekday daytime, weeknight, and weekend windows each preserve their observed energy:

$$\sum_{t \in B_b} x_t = \sum_{t \in B_b} e_t \quad \forall b, \tag{2}$$

where B_b is the set of hours in bucket b within the one-year reporting period. The first and last buckets in the reporting period are prorated when they are truncated by the analysis window. This prevents the model from forcing a full overnight charging requirement into an edge segment, such as trying to fit all overnight charging into only the 10 PM to midnight portion visible at the end of the reporting year.

The monthly peak variables bound managed hourly charging:

$$0 \le x_t \le p_m \quad \forall t \in T_m, \tag{3}$$

where T_m is the set of hours in month m. For each optimized tariff strategy, charging is prohibited[2] in the tariff windows targeted for demand avoidance:

$$x_t = 0 \quad \forall t \in F, \tag{4}$$

where F is the relevant forbidden on-peak set. The Schedule 6A smart energy management case prohibits weekday 6 PM to 10 PM charging. The Schedule 8 case uses a broader union of possible Schedule 8 on-peak windows, weekday 6 AM to 9 AM and 3 PM to 10 PM, so that the hypothetical managed profile avoids all modeled Schedule 8 demand periods. This is conservative because not all seasonal peak-demand periods are enforced at the same time, but it avoids overstating the savings available from Schedule 8 smart energy management. In transition hours near the start and end of route blocks, managed charging is also constrained not to exceed the observed charging level. This keeps the modeled profile from creating charging shoulders that would be operationally impossible because buses have already left the depot or have not yet returned.

6.1.3 Results

6.1.3.1 Observed Charging Behavior

The hourly energy reports show that most charging energy is delivered after morning pickup and after afternoon drop-off. The morning return period is slightly larger than the afternoon return period, though both are important charging windows. Figure 22 shows the observed total energy by hour of day, charger type, and charger.

The 15-minute power data provides a higher-resolution view of charging shape during the February through April 2026 overlap period. This is the period for which both the long-range hourly energy records and the shorter-history 15-minute power endpoint were available. Figure 23 shows a representative school-day segment using both the hourly energy-derived demand and the 15-minute power data. Figure 25 then shows the peak-hour comparison for each overlapping month.

The higher-resolution data also reveals a data-quality issue. In every overlapping month, the

[2] A production smart-charging system should allow charging during a demand-avoidance window if that charging is necessary to satisfy bus service requirements. In this screening model, the source district's observed duty cycle did not require charging during those periods, so the model used a simpler hard prohibition. If those prohibited hours had been necessary to preserve the modeled operational energy requirements, the optimization problem would have become infeasible.

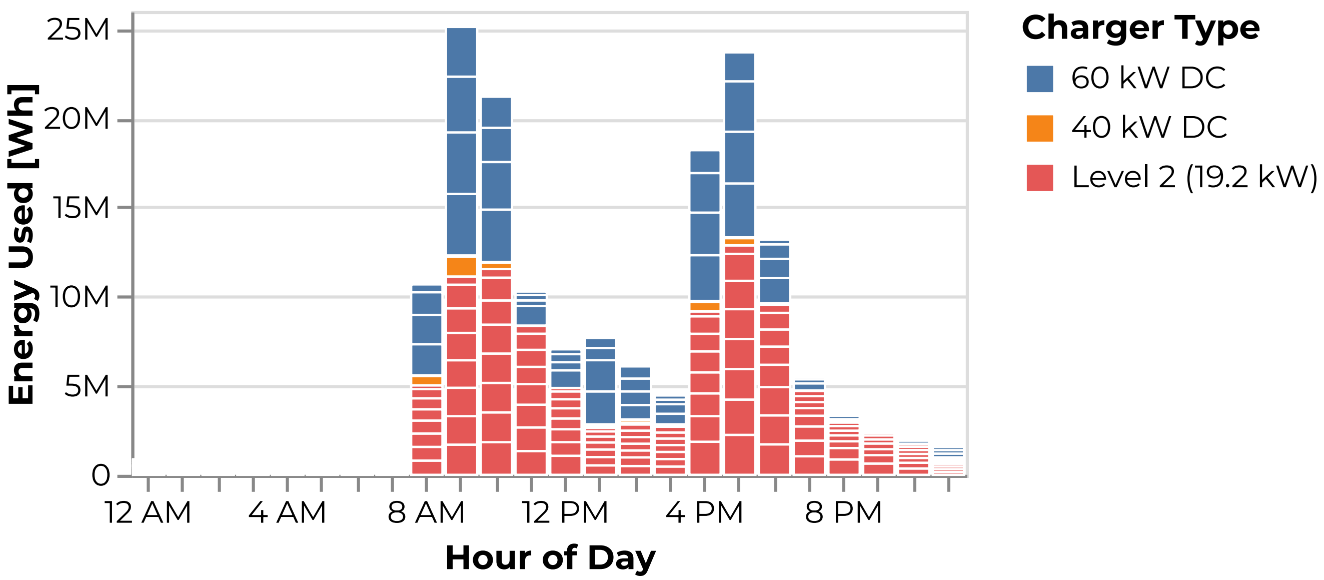

Figure 22: Observed hourly energy use by hour of day, charger type, and charger. Color indicates charger type and rated power; individual chargers remain as separate stacked segments within each charger type.

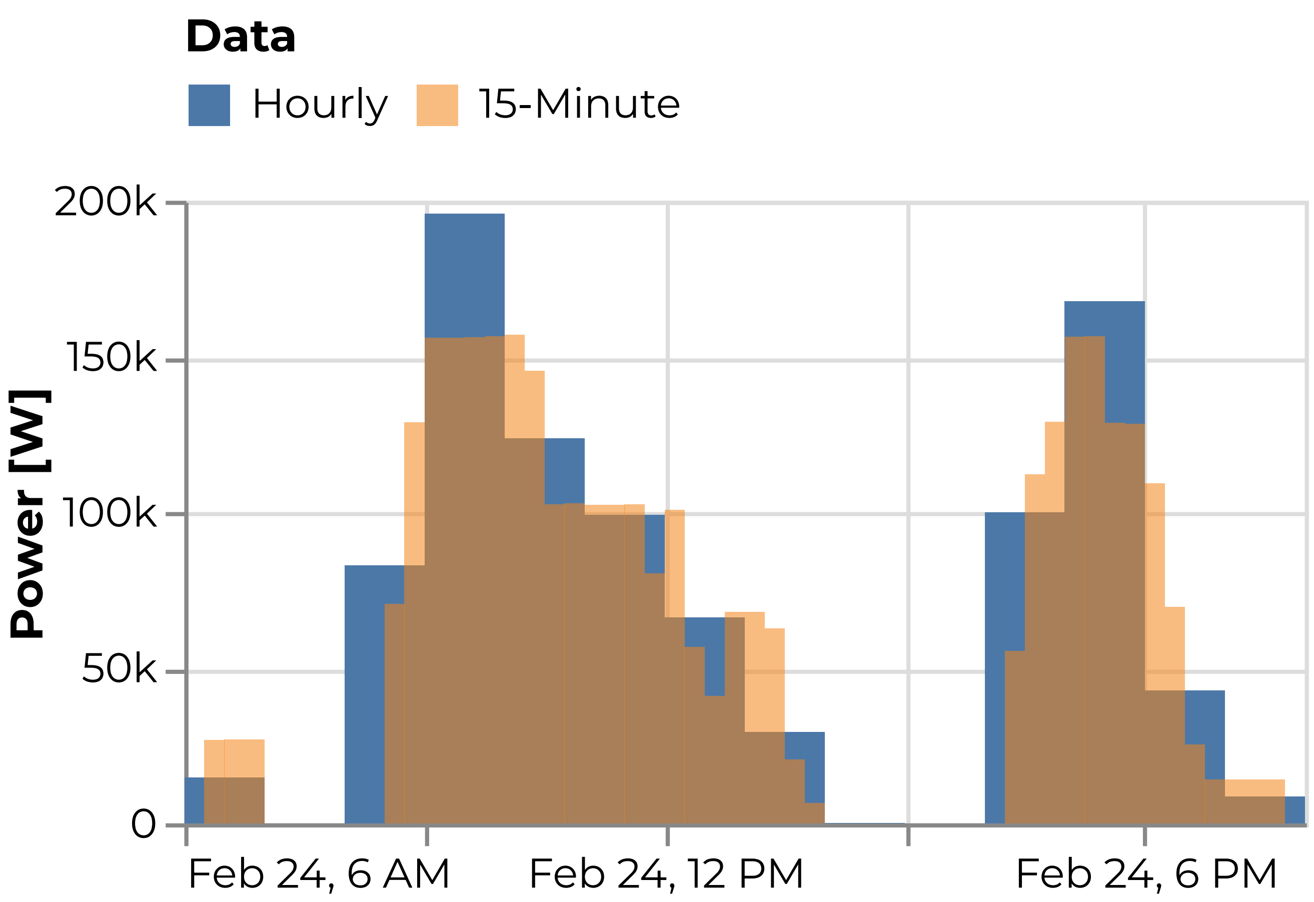

Figure 23: Observed hourly and 15-minute power comparison for a representative school-day charging segment.

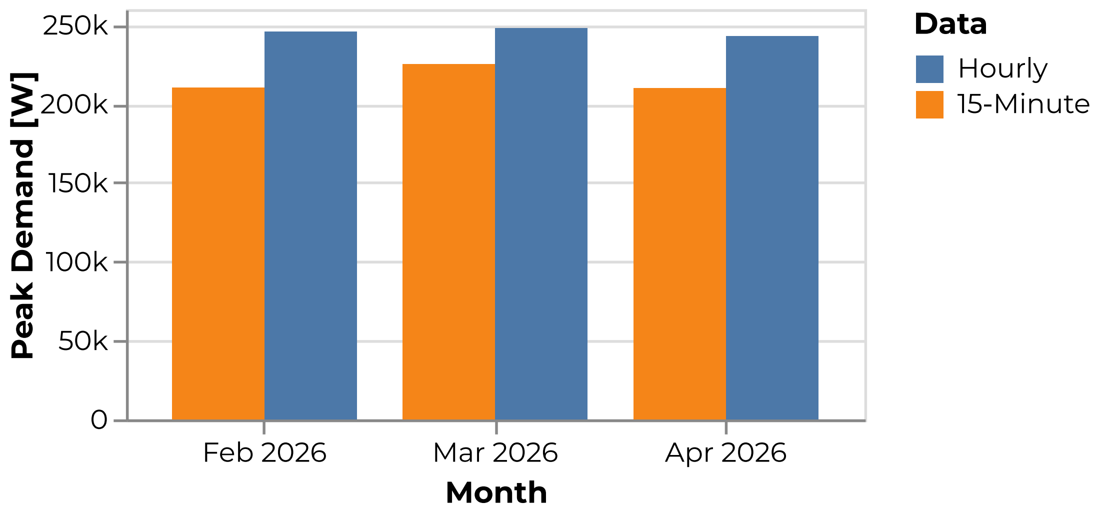

Figure 24: Observed monthly peak demand from hourly energy reports and 15-minute power data.

peak demand calculated from the hourly report exceeds the maximum 15-minute power value, as shown in Figure 24. In February 2026, the hourly peak is 246 kW while the 15-minute peak is 210 kW. In March 2026, the hourly peak is 248 kW while the 15-minute peak is 225 kW. In April 2026, the hourly peak is 243 kW while the 15-minute peak is 210 kW. A maximum 15-minute demand should not be lower than the highest hourly average if the 15-minute endpoint represents complete interval averages aligned to the same load. If the endpoint is sampled, rolling, rounded, filtered, or missing intervals, however, those endpoint behaviors could account for the discrepancy. These differences suggest the platform data should be used as a strong operational signal rather than a perfect revenue-grade meter record. For future fleet planning and performance evaluation, the best practice is to preserve both charger session-level records and 15-minute interval data from the utility meter and/or the EVSE itself.

Figure 25 compares the hour containing the monthly peak from each perspective. The patterns are plausible and useful, but the mismatch in peak magnitude and timing reinforces the need to compare the Fleetbox-derived peaks with the utility-billed monthly peak demand and, when available, Rocky Mountain Power 15-minute interval data before using the profile for billing decisions or infrastructure sizing.

6.1.3.2 Rate Comparison

With no change in charging behavior, Schedule 6A is the lowest-cost immediately actionable schedule. Figure 26 shows that Schedule 6A is below Schedule 6 in all non-summer-break months, and the annual comparison in Figure 27 shows the scale of the difference. These are simulated bills for the observed EV charging load, not actual historical bills. Schedule 8 is included to show potential tariff value if the load profile grows beyond the applicability threshold, but the district's observed charging load does not naturally meet that threshold.

The simulated annual Schedule 6 bill for the observed EV charging load is about $55,200, compared with about $36,700 under Schedule 6A. That is an annual savings opportunity of about $18,500, or 34%, before changing any charging behavior. The simulated observed-load Schedule 8 bill is about $51,600, which is about $3,500 below the current Schedule 6 baseline but about $15,000 above Schedule 6A. This makes Schedule 6A the clearest near-term actionable result of the analysis.

6.1.3.3 Smart Energy Management Results

Smart energy management creates additional savings opportunities by shifting charging away from expensive demand windows while preserving the modeled charging energy in each operational bucket, as described in Section 6.1.2.3. Figure 28 shows how the monthly bills change af-

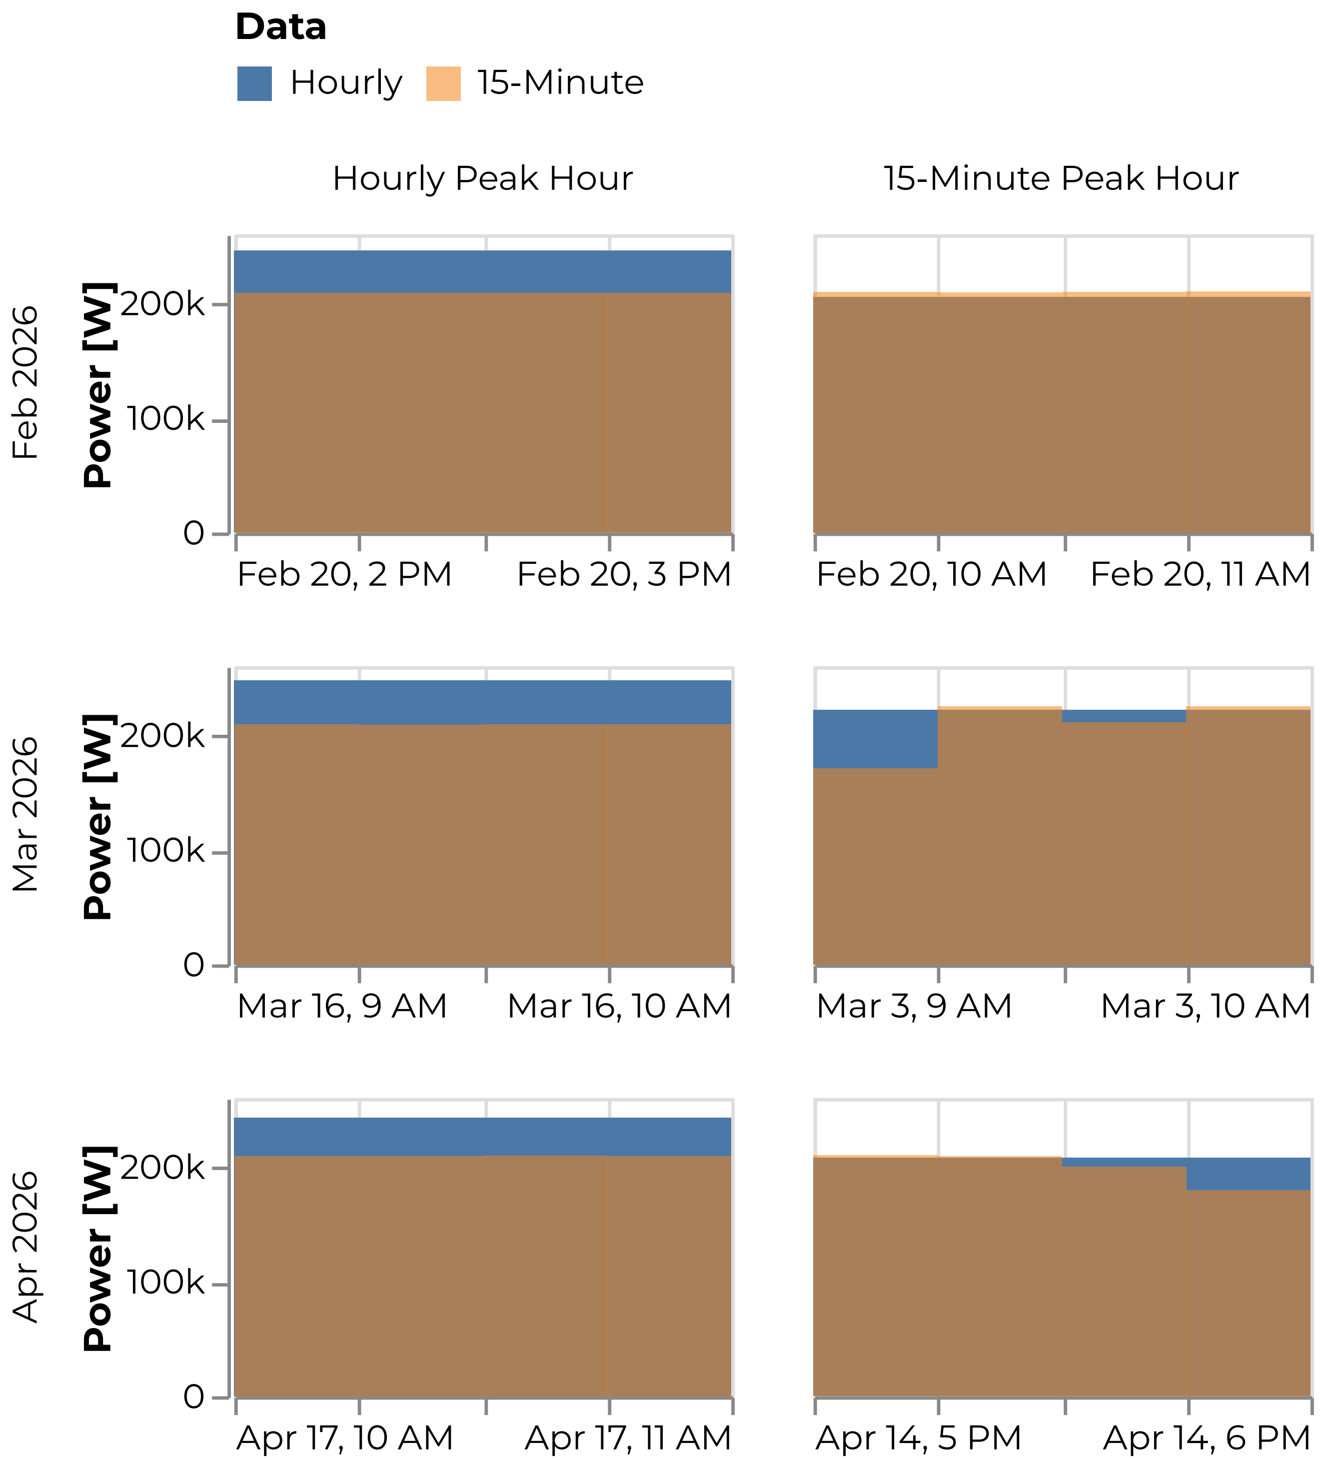

Figure 25: Detailed comparison of observed hourly and 15-minute power in the hour selected by each monthly peak-demand perspective.

ter applying the optimized charging profiles. This monthly view matters because the value of smart energy management is not only annual; it shows whether savings are persistent across the school year or concentrated in a few months. Figure 29 then summarizes the annual bill and unit-cost comparison for the managed cases. Schedule 8 is not an immediately available tariff option for

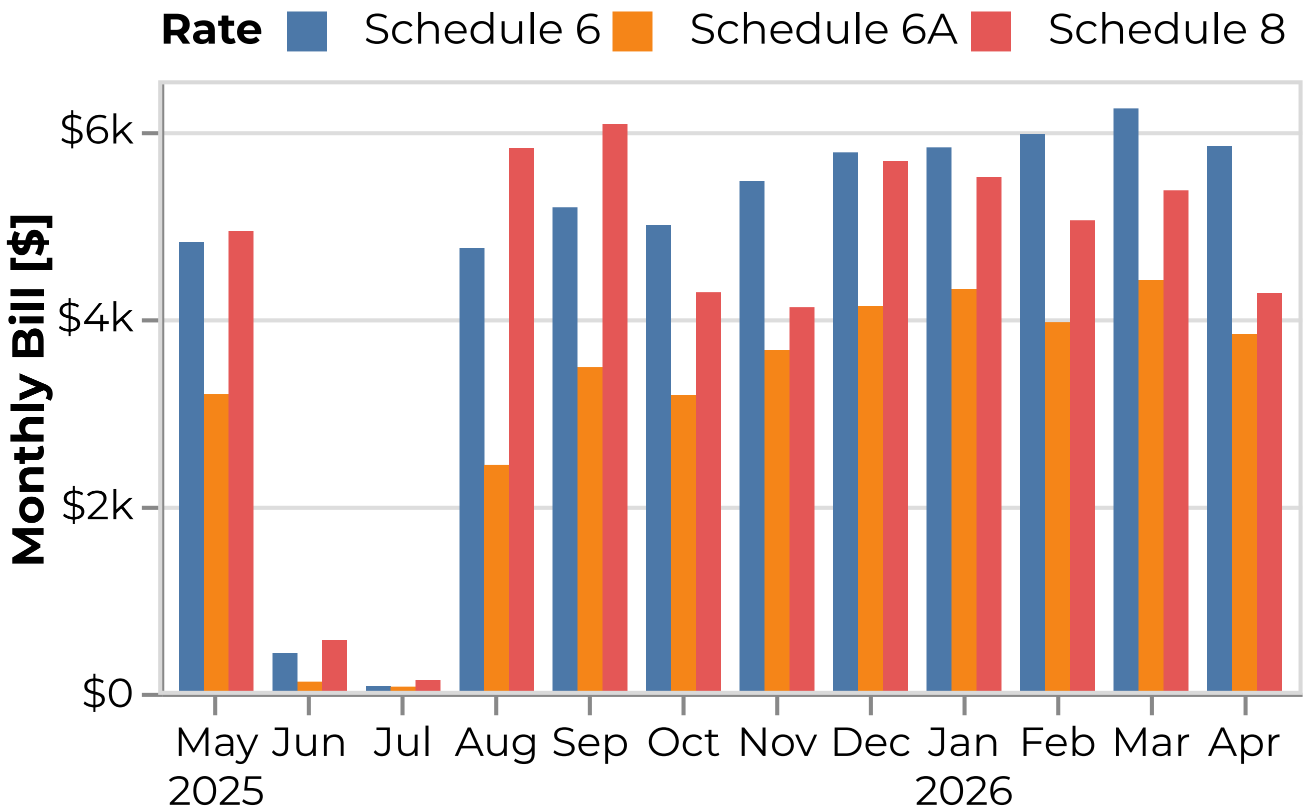

Figure 26: Simulated monthly bill estimate for the observed EV charging load by Rocky Mountain Power rate schedule.

the current observed load, but it becomes much more attractive if the district's future load is large enough for Schedule 8 to apply.

The modeled Schedule 8 annual bill shown here should not be interpreted as the expected total bill for a future Schedule 8 fleet using smart energy management. A larger fleet would use more energy, so its total bill would be higher than the case modeled for the district's current charging energy. However, the unit-cost dynamic of this growth should be considered: if fleet growth is paired with smart energy management that controls peak demand and on-peak energy, some additional energy could be served with a similar facility-charge burden. Because the marginal cost of managed off-peak energy is lower than the average cost shown for today's smaller load, a larger well-managed fleet could have an average cost per kWh below the smart Schedule 8 unit cost estimated in this analysis.

Table 6 summarizes the main decision cases. Moving from current Schedule 6 billing to Schedule 6A saves about $18,500 per year with no behavior change. Adding smart energy management on Schedule 6A increases the total savings to about $29,200 per year. The modeled smart Schedule 8 case saves about $40,400 per year relative to the current Schedule 6 baseline and about $11,200 per year relative to smart Schedule 6A, but this is a tariff-comparison result rather than a recommendation to move onto Schedule 8 immediately.

6.1.3.4 Schedule 8 Future Applicability and Rate Design

The Schedule 8 results are most useful as a future-planning signal. If the district adds enough buses and chargers for the site load to exceed 1 MW, even just once every 18 months, Schedule 8 would become the relevant large-customer tariff framework. The value of avoiding Schedule 8 peak-demand windows would then be substantial because Schedule 8 reduces both energy

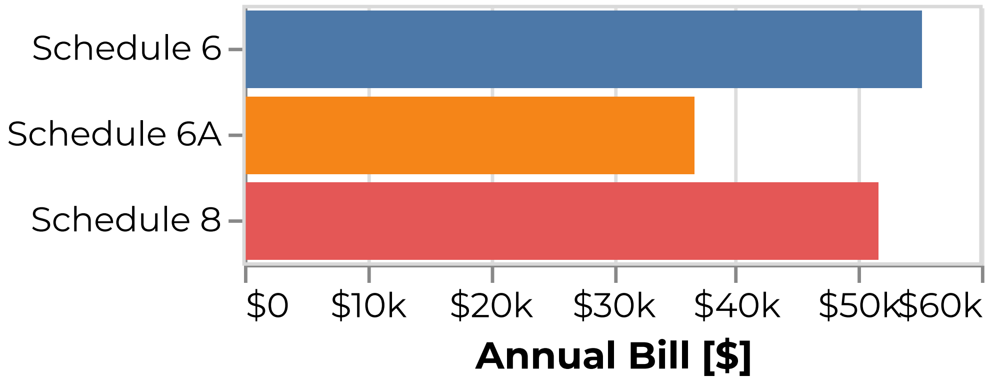

a(a) Simulated annual bill for the observed EV charging load by rate schedule.

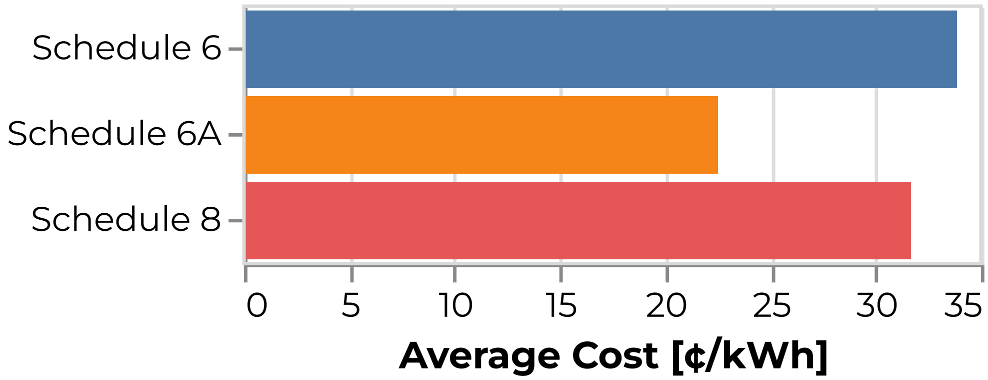

a(b) Simulated average unit cost for the observed EV charging load by rate schedule.

Figure 27: Simulated annual bill and average unit cost for the observed EV charging load by Rocky Mountain Power rate schedule.

charges and peak-demand charges when load is moved out of time-defined on-peak periods; that time-based demand-charge incentive is the key feature that separates Schedule 8 from Schedule 6 and Schedule 6A. Figure 30 compares the monthly peak-period demand and peak-period energy that would occur inside the Schedule 8 demand windows for the observed charging profile, the smart Schedule 6A profile, and the smart Schedule 8 profile.

One additional consideration that must be taken into account when switching to a rate like schedule 8 is that it has two effective demand charges. The first is a low facilities charge that applies to the the all-time peak demand for the month. The second is a much higher (~4x) peak demand charge that applies only to demand within the peak demand periods each month. Because in this analysis the district is able to eliminate all charging during the peak demand period this cost is $0. However, if there were some kind of special circumstance and charging were required during this period, even just once in a given month, the costs would be high. If just a single bus charged at 60 kW for 15 minutes during the peak demand period the peak demand charge would be roughly ~$1200, resulting in a unit price of nearly $80/kWh if that were the only such session during the month. It is therefore critical that the district commit to and be able to reliably avoid the peak demand periods before any change to schedule 8. It would likely be worthwhile for the district to evaluate public charging, for such exigent circumstances.

This comparison also points to a broader rate-design opportunity. Schedule 6A gives the district a

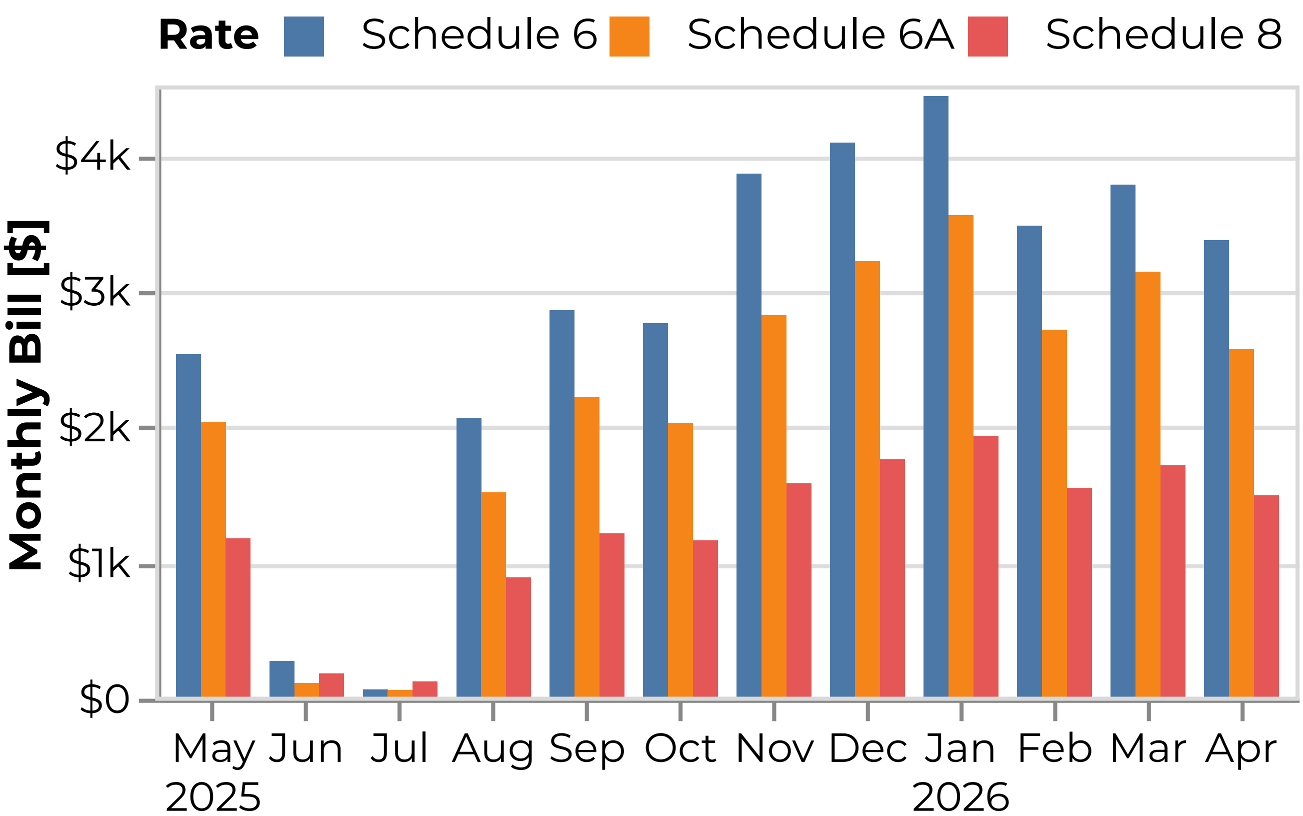

Figure 28: Monthly bill estimate after modeled smart energy management.

meaningful near-term savings opportunity, but it does not create the same strong incentive to reduce demand during Schedule 8 peak-demand windows. A future commercial time-of-use rate for smaller loads, structured more like Schedule 8 but without requiring a customer to reach 1 MW, could reduce the district's bill while also helping Rocky Mountain Power manage peak-period load. Such a rate would better align customer savings with utility-system benefits for school-bus fleets and other flexible commercial loads that are too small for Schedule 8.

The average weekday load profiles in Figure 31 show how the optimized charging strategies reshape the day. The Schedule 6A strategy primarily moves energy away from weekday evening hours. The Schedule 8 strategy is more aggressive because it avoids both morning and broader afternoon/evening demand windows.

6.1.4 Conclusion and Recommendations

The analysis found three decision-relevant results. First, the observed charging profile is strongly shaped by morning and afternoon route-return windows. Second, Schedule 6A materially reduces annual cost even without changing charging behavior. Third, smart energy management can create additional savings if it is implemented without compromising route readiness.

The first recommendation is to immediately prioritize switching from Schedule 6 to Schedule 6A. The modeled annual savings are about $18,500, or 34%, with no behavior changes. This is a rate-selection decision rather than an operations change, so it should be pursued before more complex control strategies. Because the district's current charging load is below the Schedule 8 threshold, Schedule 6 versus Schedule 6A appears to be the practical near-term rate choice. The district should still confirm the tariff change

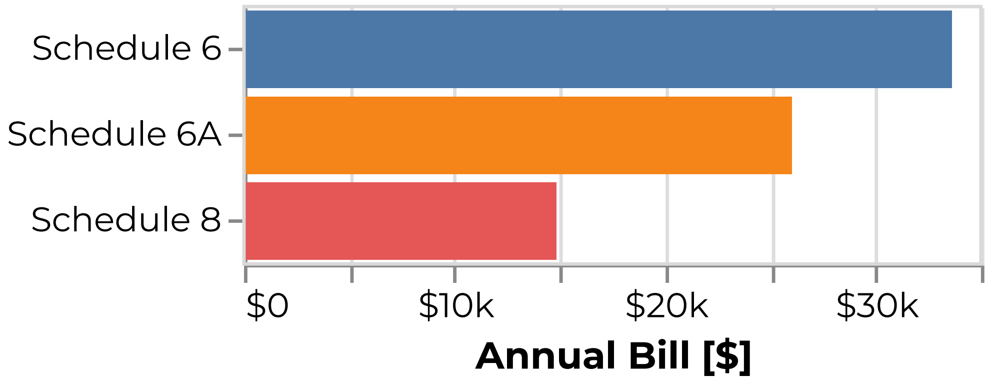

c(a) Annual bill after modeled smart energy management.

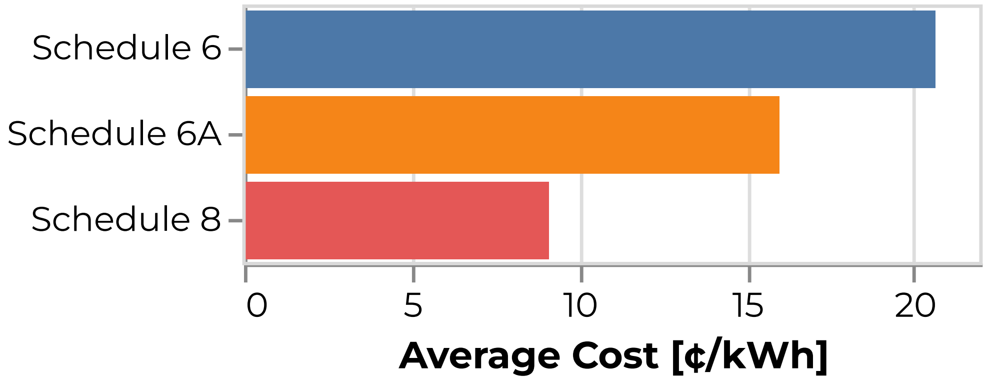

c(b) Average unit cost after modeled smart energy management.

Figure 29: Annual bill and average unit cost after modeled smart energy management[3].

Table 6: Annual cost and savings summary for the observed and managed charging cases[4]; savings are relative to current observed charging on Schedule 6, and Schedule 8 values are modeled tariff comparisons.

Option	Annual bill	Savings [$]	Savings [%]	Context
Current Schedule 6	$55,183	$0	0%	Current baseline
Observed Schedule 6A	$36,661	$18,523	34%	Rate switch only
Smart Schedule 6A	$26,016	$29,167	53%	$10,645 below observed Schedule 6A
Smart Schedule 8	$14,820	$40,364	73%	$11,196 below smart Schedule 6A

with Rocky Mountain Power and review any applicable service agreements before switching.

The second recommendation is to begin experimenting with smart energy management during the summer period. Summer provides a lower-risk testing window for delayed charging, charger response, operational monitoring, and route-readiness checks before staged fall-semester deployment. If testing begins while the site remains on Schedule 6, the district should avoid unnecessary tests that concentrate charging into a new monthly demand peak, because Schedule 6 bills the full monthly demand charge. Sched-

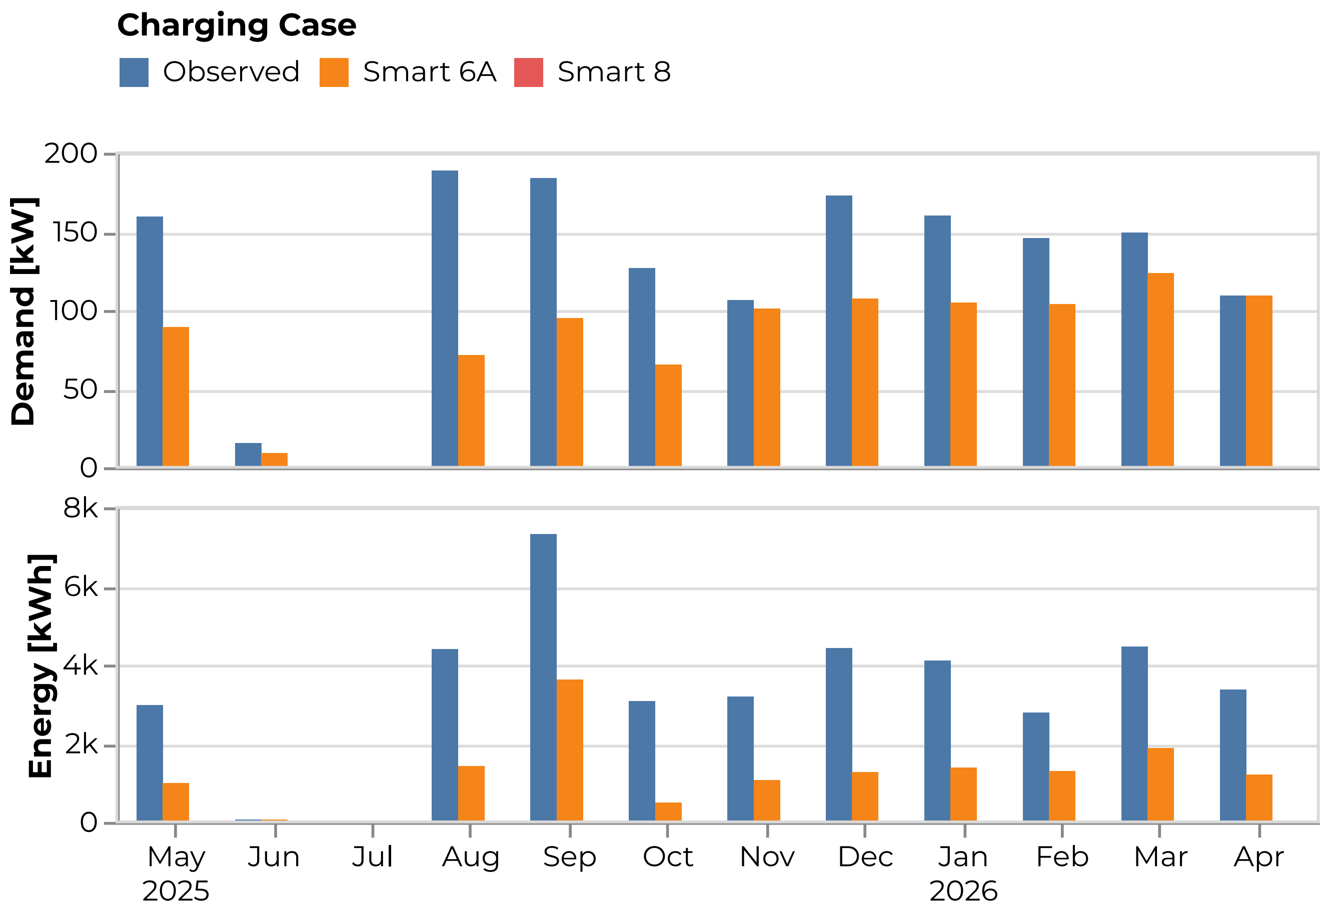

Figure 30: Monthly peak-period demand and energy inside Rocky Mountain Power Schedule 8 demand windows for the observed charging profile and modeled smart energy management profiles. Smart Schedule 8 bars are zero, not missing.

ule 6A provides a safer setting for these tests because its peak-related charge is spread over energy use rather than charged as a full monthly demand charge. The first phase should focus, where possible, on buses whose overnight state of charge would still support the next morning route if managed charging failed to resume after an evening demand-avoidance period[5]. Once the team gains confidence that smart energy management is reliable and does not create route-readiness issues, it can be expanded to more buses and more nights.

Operational smart energy management should be treated as a follow-on implementation project rather than a setting that can be fully specified from the available Fleetbox exports alone. The control strategy should be configured and tested so that three requirements are satisfied: buses receive the energy needed for their assigned routes, the site peak is held below a month-specific demand target where practical, and charging energy is limited to off-peak periods unless charging during a peak period is necessary for operational readiness. The source district should work with ASPIRE and the relevant charging-platform, charger, or energy-management vendors to determine what controls are available in the existing system, how those controls should

[5]This starting point reduces the operational consequence of a control or communications failure. If delayed overnight charging does not restart as expected, the affected bus should still be able to complete the morning route, and staff can diagnose the problem during normal working hours before the afternoon route window. Daytime managed charging can still protect afternoon readiness because drivers, transportation staff, and other support personnel are on site to respond if a charger or control system does not behave as expected.

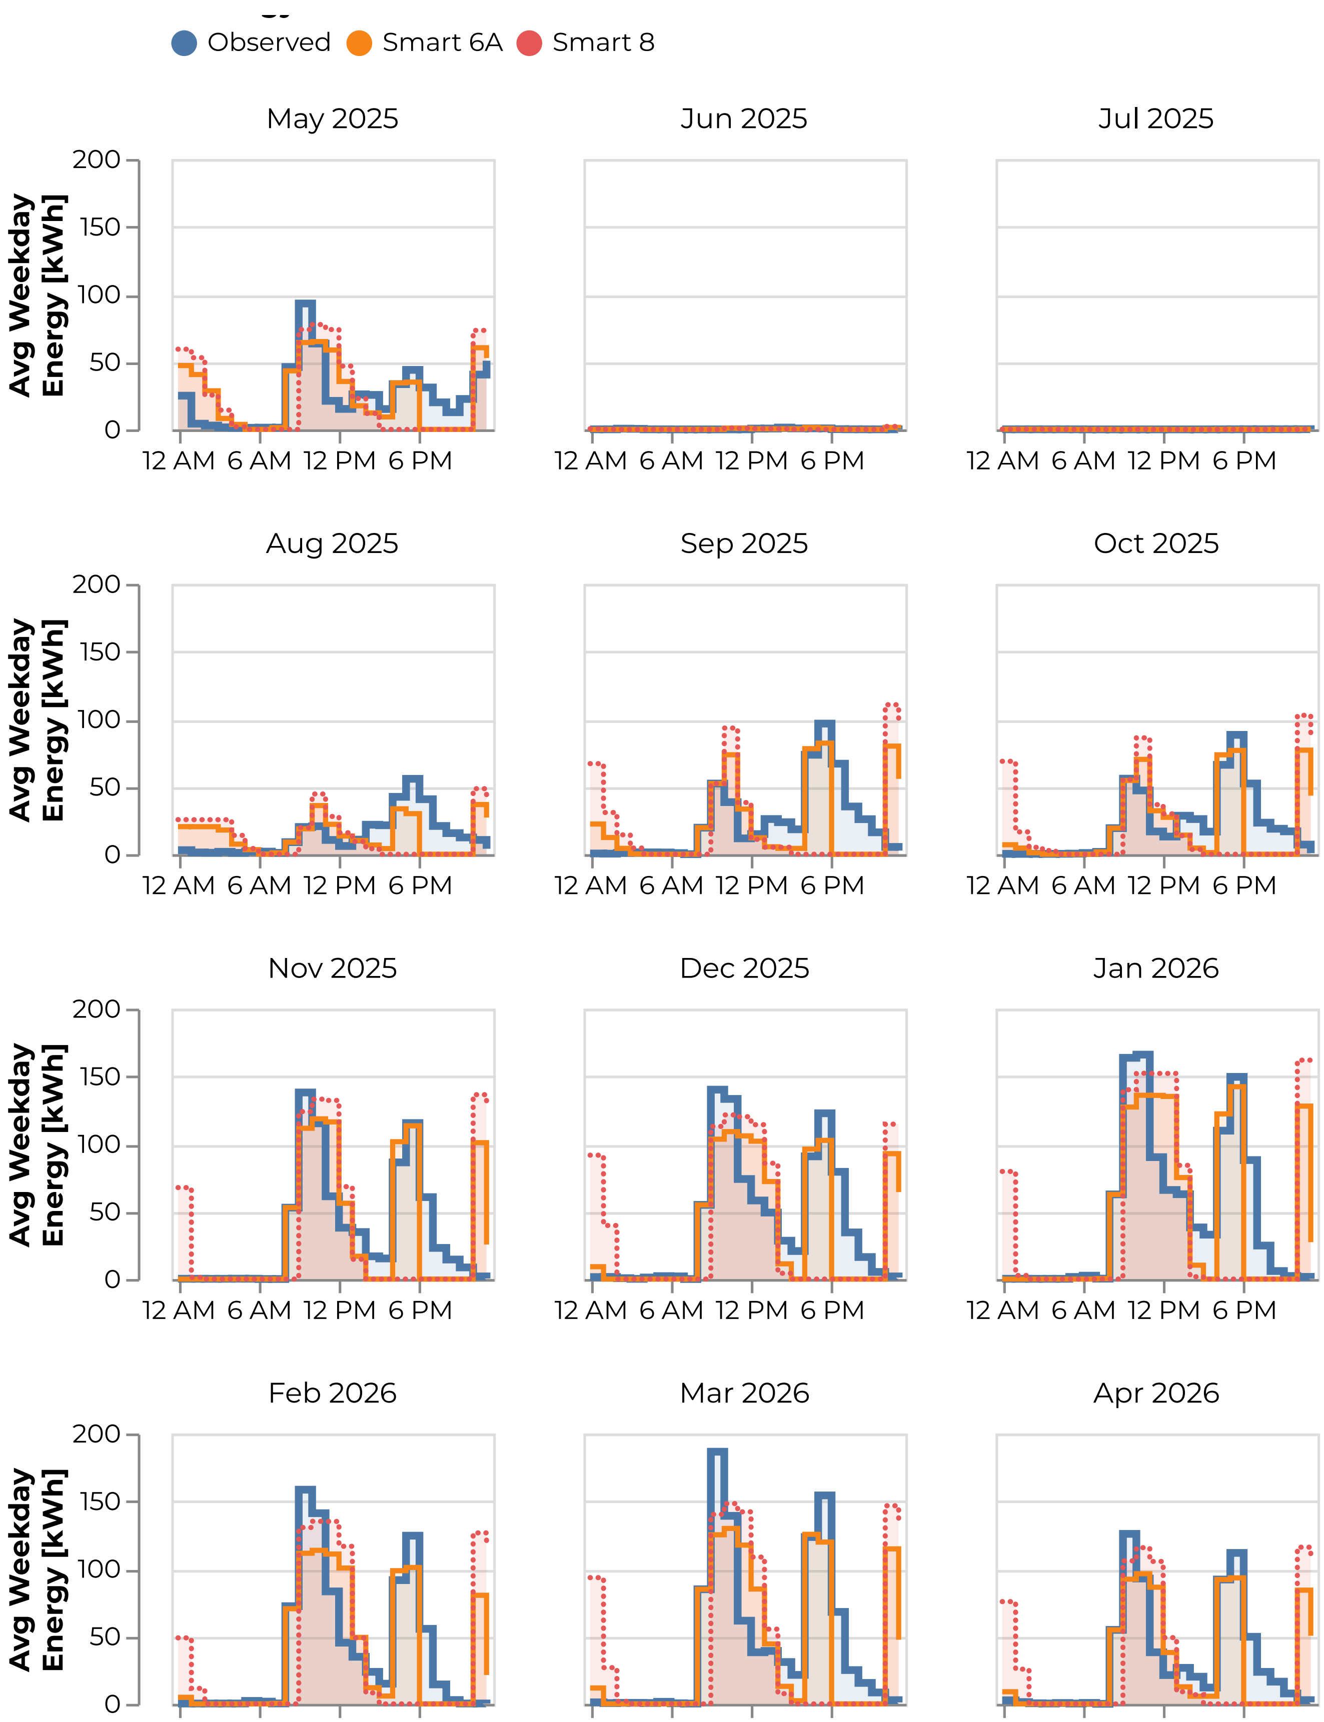

Figure 31: Average weekday charging profile by month for observed charging, smart Schedule 6A management, and smart Schedule 8 management.

be configured, and what monitoring is needed before relying on the strategy during normal operations.

Future smart energy management work would be stronger if the district can obtain and record both utility-side and charger-side data. On the utility side, monthly billing peak demand should be compared with the Fleetbox-derived peak estimates, and Rocky Mountain Power 15-minute interval or load-profile data should be requested when a more detailed operating profile is needed. On the charger side, the district should ask its existing charging-platform vendor, and any future vendors it evaluates, for accessible session-level records and higher-resolution power or energy data. Those records would make it easier to diagnose data discrepancies, evaluate managed-charging performance, and refine future control strategies.

The third recommendation is to keep Schedule 8 in view as a planning issue for future fleet expansion, not as a near-term implementation step. The district reported that Rocky Mountain Power identified a 1,200 kVA interconnect, which suggests that the utility-side service may be large enough for a future 1 MW load without a utility service upgrade. The current charging equipment, however, provides less than 500 kW of available charging capacity, and the observed charging load does not naturally reach 1 MW. Practically speaking, the cleanest path to Schedule 8 applicability would be a larger fleet and more charging equipment. As expansion plans develop, the district should review its Rocky Mountain Power service agreement, interconnection assumptions, and tariff requirements so it understands when Schedule 8 could become available and/or mandatory and what operating constraints would apply.

A broader recommendation is for Rocky Mountain Power and other utilities to consider a Schedule 8-like option for smaller flexible commercial loads. Schedule 8 provides a strong incentive to align EV charging behavior with grid needs and constraints because it rewards customers for moving both energy and peak demand away from high-cost periods. The peak-demand component is especially important: unlike Schedule 6 and Schedule 6A, Schedule 8 makes on-peak demand reductions directly valuable through time-based peak-demand charges. If that incentive structure were available below the 1 MW threshold, rural school-bus fleets with similar operating characteristics could lower their bills while also reducing utility-system peak stress. That kind of rate design could create benefits on both sides: a lower-cost path for fleet electrification and a more grid-supportive charging profile for the utility.