

The DC Corridor Playbook

Highway charging where throughput is the whole business

Corridor sites are a reliability business first and a pricing business second. This is how they are run.

FOR Operators building or running highway and forecourt sites
UPDATED 12 July 2026

Contents

01 Corridor economics

Derated power, session length, bays per site and the throughput arithmetic that decides payback.

02 Reliability as the product

What a broken corridor charger actually costs, and the operational habits that prevent it.

03 Peaks and queueing

Festival weekends, holiday traffic, and pricing the bay rather than only the electrons.

04 Forecourt partnerships

Working with fuel retail — staff who are not charging specialists, existing POS, and scoped access.

05 Grid and siting

High-power connections, sanctioned load, and why an existing sanction is worth more than a better location.

06 Being findable

Live availability in the apps drivers already use, and what roaming is worth on a corridor specifically.

WHAT YOU GET

- Sizing a corridor site for throughput rather than for connector count
- Why a stranded driver costs more than any single repair
- Idle fees and queueing on a site with genuine peaks
- Working with fuel retail hosts and their existing operations

Corridor economics

A corridor site is a throughput business. Nobody is dwelling, nobody is shopping, and the only thing that produces revenue is vehicles cycling through bays. That makes the arithmetic unusually clean and unusually unforgiving.

Start from derated power

A 60kW charger does not deliver 60kW for an hour. Vehicles taper as the battery fills, the unit derates when hot, and there is dead time between sessions while one driver unplugs and the next arrives.

Realistic delivered energy is somewhere between 55 and 70 percent of nameplate multiplied by connected time, and the figure is lower in summer. Every model built on nameplate is wrong by roughly a third.

The throughput chain

STEP	EXAMPLE
Nameplate per connector	60 kW
Realistic average delivery	38 kW
Mean connected time per session	25 minutes
Energy per session	~16 kWh
Dead time between sessions	8 minutes
Sessions per connector per busy hour	~1.8
Busy hours per day at a good site	6-8

Multiply through and you have units per connector per day, which multiplied by contribution per unit is the site’s daily contribution. Against fixed monthly cost, that gives payback — and it is a far more honest calculation than a utilisation percentage.

Dead time is a real lever

EIGHT MINUTES OF DEAD TIME IS A FIFTH OF THE SESSION

On a 25-minute session, dead time between vehicles is a large share of the cycle. Bay layout that lets one car leave while the next approaches, cables long enough to reach either side, and clear signage do more for throughput than a power upgrade.

How many bays

Fewer than instinct suggests, and never one. One bay means a single fault takes the site to zero and a single occupied connector turns arrivals away. Two is the practical minimum for a corridor site; four is where queueing stops being the constraint at ordinary traffic.

Beyond that, additional bays serve peaks only, and the question becomes whether the peak is frequent enough to justify capital that idles the rest of the week.

What kills corridor returns

- 01 Downtime, which is throughput permanently lost — the session does not happen later.
- 02 Oversizing the sanction, which adds a fixed monthly cost the throughput has to carry.
- 03 Overestimating delivered power, which inflates every derived figure by a third.
- 04 Underestimating dead time, which is the difference between two sessions an hour and one.
- 05 A location that is convenient on a map and awkward in a vehicle.

Reliability as the product

On a destination site, a broken charger is an inconvenience — the driver shops and comes back tomorrow. On a corridor, it is a driver who cannot complete their journey. The cost is not the session; it is the person who tells everyone, and the drivers who stop trusting the route.

The asymmetry

A corridor site that works is invisible. A corridor site that fails once is remembered, reviewed, and routed around for months by people who never experienced the failure themselves. There is no equivalent upside to working well, which means reliability spending on a corridor is loss-avoidance rather than revenue generation — and should be budgeted as such.

The habits that prevent it

- 01 Never migrate, upgrade or experiment on a corridor site first. It goes last in every rollout, every firmware cohort, every platform change.
- 02 Redundancy at the bay level. Two smaller units beat one larger one, because a fault takes half the site rather than all of it.
- 03 A response time commitment for corridor sites that is shorter than for the rest of the network, and a technician who can actually meet it.
- 04 Proactive alerting on the zero-energy case, not only on faults. A unit reporting Available and delivering nothing is the worst failure mode here.
- 05 A daily glance at corridor sites specifically, separate from the network dashboard.

THE ALERT THAT MATTERS MOST ON A CORRIDOR

Not "charger offline" — everyone has that. It is "charger online, available, and no session in N hours at a site that normally has them". A corridor site with an unexplained quiet afternoon is either a fault the protocol cannot express or a blocked bay, and both need someone to look.

Spares and access

A corridor site is usually far from your depot, which is exactly why it needs local provision. A small spares kit held at the site — a connector assembly, a router, a fuse set — turns a day-long dispatch into a call with the forecourt manager.

And access has to work at the hours the site is busy. A technician who cannot get in at 22:00 because the site staff have left is a technician who arrives at 09:00 to a charger that has been down all night.

When it does fail

- 01 Mark it unavailable immediately, so nobody diverts to it. A charger shown as available and broken is worse than one shown as broken.
- 02 Push the status to roaming partners and map providers, not just your own app.
- 03 If the site has other working bays, say so — a partial outage should not read as a closed site.
- 04 Tell drivers who were en route, where you can.

The first point is the one most often missed and the cheapest to fix. Most of the damage from a corridor failure is done to the drivers who arrive after it, not the one who found it.

Peaks and queueing

Corridor demand is not evenly distributed. It is concentrated into weekends, holidays and festival travel, when a site that is comfortable all week develops a queue. Managing that peak is a pricing and operations problem rather than a capital one, because building for the peak means idling for the rest of the year.

Know your peaks in advance

Indian corridor traffic follows a predictable calendar: long weekends, major festivals, school holidays, and the days on either side of them. These are knowable months ahead, which makes them plannable in a way that a random surge is not.

- Staff differently on those days, or at least have someone reachable.
- Do not schedule maintenance anywhere near them.
- Make sure spares are on site beforehand.
- Expect thermal derating on the busiest afternoons, and plan throughput accordingly rather than being surprised.

Pricing the bay

At peak, the scarce resource stops being energy and becomes the connector-hour. Pricing only energy at that point rewards the driver who charges to 100 percent — slowly, in the taper, occupying a bay for twenty minutes to add very little — at everyone else's expense.

THE TAPER IS WHERE QUEUES FORM

The last twenty percent of a battery takes a disproportionate share of the time and delivers a small share of the energy. A pricing structure that makes that expensive moves drivers on at the point where their marginal benefit is smallest, which is exactly where you want the behaviour change.

Instruments, in order of preference

- 01 A strict idle fee with a short grace period. Free the bay the moment charging finishes.
- 02 A time component alongside the per-kWh rate, so the connector-hour is priced.
- 03 Time-of-day pricing where the peak is predictable by hour rather than by day.
- 04 A session cap at very constrained sites, disclosed clearly — which is heavy-handed and occasionally correct.

What not to do is surge pricing on a corridor. A driver mid-journey has no alternative, so a surcharge is a toll rather than a signal, and it will be remembered as one.

Managing the queue itself

When a queue forms anyway, the things that help are unglamorous: clear signage about where to wait, live availability published so drivers can see the wait before exiting the highway, and — if the site has staff — somebody keeping order at the busiest hours.

A reservation system sounds like the answer and generally is not, because a reserved bay sitting empty while a queue waits is worse than no reservation at all.

Forecourt partnerships

Fuel retailers are the natural hosts for corridor charging: they have the land, the connections, the amenities and the traffic. They also have an existing operation with its own systems, its own staff and its own commercial expectations, and a charging operator who ignores that will be a difficult partner.

What they bring

- A location already chosen for passing traffic, which removes most of the site-selection risk.
- An electrical connection, often with headroom, and a relationship with the DisCom.
- Amenities: a shop, a washroom, food. Drivers need somewhere to be for twenty minutes.
- Staff on site, which is genuinely valuable for a corridor charger.
- Lighting, cameras and a security presence overnight.

What they are worried about

- 01 Forecourt space. Every bay you take is space that had a use, and they will want to know it is a good trade.
- 02 Their staff being asked to support something they do not understand.
- 03 Safety and compliance on a site that already has strict rules for good reasons.
- 04 Being tied to one charging operator if the market shifts.

- 05 Whether charging customers actually spend in the shop, which is the question they most want answered.

ON THE SHOP-SPEND QUESTION, BE HONEST

A charging driver dwelling for twenty minutes does spend more than a fuel customer dwelling for three. The effect is real, modest, and hard to attribute. Overclaiming it is the fastest way to lose credibility with a retailer who has heard the pitch from everyone.

Scoped access for site staff

Forecourt staff should be able to help a stuck driver without being able to break anything. That means a login limited to their own site, with operate-but-not-manage permissions: start, stop, reset, set unavailable, and nothing else.

Give them a one-page card: what the indicator means, the two things to try, and when to call you. Staff who can resolve the common cases save you dispatches and give the driver a better experience than any remote support could.

Payment and the existing POS

Retailers frequently want charging to appear on their till, so a driver pays for a coffee and a charge together. That is technically achievable through the API and it is a real integration project rather than a configuration setting.

It is worth doing where the retailer is a serious partner with many sites, and worth declining politely for a single forecourt — the effort does not scale down.

Compliance on a fuel site

A forecourt has hazardous-area classifications, separation distances and rules that predate electric vehicles by decades. Involve the retailer's own safety function early rather than presenting a design. They are not obstacles; they are the people who know which parts of the site you cannot use, and finding that out at the design stage is much cheaper than at installation.

Grid and siting

Corridor charging needs real power, and real power is the slowest thing to obtain in this business. That single fact should dominate site selection far more than it usually does.

The trade nobody makes explicitly

A SLIGHTLY WORSE LOCATION WITH POWER BEATS A BETTER ONE WITHOUT

A site with an existing high-capacity connection can be live in weeks. A better site needing a new sanction and possibly a transformer can be live in a year, maybe. In a market where position matters, twelve months is usually worth more than a marginally better location.

What to establish before anything else

- 01 Sanctioned load and recorded maximum demand. The gap is free headroom.
- 02 What the existing connection physically is — voltage level, transformer capacity, spare ways.
- 03 Whether the DisCom has capacity in the local network, which is a different question from whether you can pay for a sanction.
- 04 How demand charge is assessed here, and whether an EV concession applies.
- 05 Realistic timelines for an increase, from someone who has actually done one recently in that circle.

Sizing

Size against coincident peak, not the sum of connector ratings. Corridor sites do have less diversity than destination sites — arrivals cluster — but they still never draw nameplate simultaneously, because vehicles taper at different points and dead time is real.

Load management matters here too, and differently: on a corridor the constraint is throughput, so a ceiling that slows every vehicle equally is worse than one that lets an arriving vehicle take full power while a tapering one gives some back.

Physical siting on the plot

- Approach and exit that work for a car with a trailer or a light commercial vehicle, not just a hatchback.
- Cable reach to either side of a vehicle, because charge ports are on different sides.
- Bollards, without which a reversing vehicle eventually finds the charger.
- Shade, which is worth real throughput on a summer afternoon through avoided derating.
- Distance from the pumps that satisfies the retailer's safety rules, established before the layout is drawn.
- A cable route that does not cross the forecourt's busiest path.

Planning phase two at phase one

Corridor sites that work get more bays. Spare ways in the board, conduit capacity in the trench, physical space at the panel and a feeder sized for the eventual ceiling cost very little during the original civil work.

Trenching a forecourt twice is worse than trenching a car park twice, because the second time you are closing lanes at a working fuel station.

Being findable

A destination charger is found by someone who was coming anyway. A corridor charger has to be found by someone deciding, at speed, whether to leave the highway. Discoverability is not marketing here — it is the difference between a site working and a site not.

Where drivers actually look

- 01 The vehicle's own navigation, on newer cars.
- 02 A general map application, which is where most people start.
- 03 A charging aggregator app.
- 04 Their own provider's app, if they have an account somewhere.
- 05 Your app, last, and only if they already have it.

An operator relying on the fifth is relying on the smallest of the five. On a corridor, where the driver is often outside their usual area, it is smaller still.

Live status, not just a pin

A PIN WITHOUT STATUS COSTS YOU THE DRIVER TWICE

They divert, find it occupied or faulted, and both the wasted trip and the site are remembered. Published status that says "one of two available" lets them decide before exiting, and a driver who chooses not to come is a far better outcome than one who comes and cannot charge.

What roaming is worth here specifically

More than anywhere else in a network. A corridor driver is by definition away from home, likely to hold an account with a provider that is not you, and unwilling to install an app at the roadside. Roaming converts them from a lost session into a paying one.

And the substitution risk that argues against roaming at busy urban sites is weaker here outside peaks — a corridor site at 14:00 on a Tuesday has capacity that would otherwise earn nothing.

The details that matter at speed

- Accurate coordinates that lead to the entrance, not the centre of the plot.
- A name a driver recognises — the fuel brand and the highway, not an internal site code.
- Access notes: which entrance, whether it is 24 hours, whether a barrier is involved.
- Photographs of the actual bays, which do more for confidence than any description.
- Connector types stated plainly, because a driver arriving at the wrong connector is a wasted diversion.

Signage from the road

The least technical item and among the most effective. A driver at 80km/h needs to know a charger is there before the turning, not at it. Where the retailer allows signage, it is a cheaper source of sessions than almost anything digital, and it

works for the driver who never opened an app at all.

About this document

Written by the team building zevOS, the charge point management platform. We publish this kind of thing because operators who understand utilisation, demand charges and settlement make better decisions — including about us.

Corrections and disagreements are genuinely welcome. If something here is wrong, or true in most states but not yours, write to support@zevos.ai and we will fix it in the next revision.

White Stripe Innovations Private Limited. Last updated 12 July 2026. This document is general information about running a charging business, not legal, tax or financial advice.