The Uptime Playbook

Measuring, defending and improving charger availability

Why the industry's uptime numbers are not comparable, how to compute one that is, and the operational habits that move it.

FOR Operations leads and field-service managers

UPDATED 10 June 2026

Contents

WHAT YOU GET

- A defensible definition of availability, measured at the connector
- The fault triage ladder that closes incidents without a site visit
- What to stock at the depot and why the warehouse is the wrong place for it
- An SLA template that separates your obligations from the host’s

CHAPTER 01

Three definitions of uptime

Charging operators quote uptime figures in the high nineties. Drivers report that roughly one charger in five does not work. Both are being honest, and the gap between them is definitional. Until an operator picks a definition and states it, an uptime number is a number about nothing.

Connectivity

The charger is reachable: heartbeats arrive, the WebSocket is up, the platform can talk to it. This is the easiest to measure and the least meaningful, because a charger can be perfectly reachable and completely unable to deliver energy.

It is also the number most often quoted, which is why nobody believes the industry's figures. If your reported uptime is connectivity, say so, and expect a sophisticated counterparty to discount it entirely.

Availability

The connector was in a state where a driver could have started a session. Derived from the connector state timeline: Available and Occupied both count — a connector in use is working — while Faulted and Unavailable do not.

This is the right default. It is what a driver experiences, it is derivable from data every OCPP platform already has, and it is what a counterparty will insist on once they think about it for a minute.

Serviceability

A driver arriving at this connector could actually have charged. Stricter than availability, because it excludes the failures the protocol cannot express: a bay blocked by a petrol car, a cable stolen overnight while the unit still reports Available, a payment path that is down, a unit derating so hard it delivers almost nothing.

THE GAP BETWEEN THE SECOND AND THE THIRD IS YOUR REPUTATION

Availability is what you can defend contractually. Serviceability is what gets written about. An operator who only measures the first will keep being surprised by reviews that contradict their dashboard.

Picking one, and saying which

AUDIENCE	QUOTE	BECAUSE
Host or tender	Availability, per connector, with exclusions	It is defensible and evidenced
Internal operations	Serviceability	It is the only one that finds the real problems
Nobody, ever	Connectivity	It is not a measure of the product

The rest of this playbook is about computing the second so that it survives scrutiny, and about closing the gap to the third.

CHAPTER 02

Computing a number you can defend

A defensible availability figure is one where somebody who disagrees with it can be shown exactly how it was produced. That requires the underlying state history to be retained, the arithmetic to be stated, and the exclusions to be visible rather than silently applied.

The raw material

Every connector emits state changes: Available, Preparing, Charging, SuspendedEV, SuspendedEVSE, Finishing, Reserved, Unavailable, Faulted. Stored with timestamps, those form a timeline, and availability is a straightforward integration over it.

- Counts as available: Available, Preparing, Charging, SuspendedEV, SuspendedEVSE, Finishing, Reserved.
- Counts as unavailable: Faulted, Unavailable.
- Requires a decision: periods where the charger was unreachable and sent nothing at all.

The offline problem

A charger that stops communicating has not necessarily stopped working. With a local authorisation list it may serve sessions perfectly well and report them on reconnect. Treating every offline minute as downtime understates your performance; treating it as uptime overstates it.

The defensible treatment is to decide by evidence rather than by policy: if the charger later reports transactions that occurred during the offline window, it was available. If it reports nothing, treat the window as unavailable. Both are checkable after the fact, which is what makes the number defensible.

CARRY THE LAST KNOWN STATE FORWARD, BUT CAP IT

A connector that was Available when it went silent was probably still available a minute later, and probably not still available two days later. Carrying the last state forward with a cap — an hour is a reasonable default — is honest; carrying it indefinitely is how a broken charger stays at 100 percent for a fortnight.

Excluding maintenance

Planned maintenance is legitimately excluded, and the exclusion is legitimately abused. Three rules keep it honest.

01 It must be declared before it starts. A window created retrospectively is not maintenance, it is editing history.

02 It must be bounded. An open-ended window is an off switch for the metric.

03 It must appear in the report as an exclusion, with its duration, rather than silently vanishing from the denominator.

The arithmetic

For each connector: available seconds divided by (operating seconds minus excluded seconds). Aggregate across connectors by summing seconds, not by averaging percentages — a site with one connector at 99 percent and one at 50 percent is not 74.5 percent available, and averaging percentages will eventually flatter you in a way you cannot defend.

What a report must contain

FIELD	WHY IT IS THERE
Period, explicitly	So it can be reproduced
Percentage per connector	Because per-station hides the failure
Total operating seconds	So the denominator is visible
Excluded seconds, itemised	So exclusions can be challenged
The outage list	Start, end, cause — this is the evidence

The outage list is what turns a percentage into an argument you can win. It is also the part most platforms cannot produce, so confirm it exists before committing to an SLA that depends on it.

CHAPTER 03

The triage ladder

The marginal cost of running a charging network is field dispatches. So the operational question is never "how do we fix this charger" but "can we fix it from here" — and the answer improves dramatically when triage is a fixed sequence rather than an individual's judgement.

Seven rungs, always in this order. The ordering is not arbitrary: each rung eliminates more causes than the one below it, and the destructive steps are last.

Rung 1 — read the connector state

Available, Occupied, Faulted, Unavailable. A unit reporting Unavailable was usually set that way by a person, and the fix is a conversation rather than a repair. A unit reporting Available that a driver says does not work is a serviceability problem — a blocked bay, a cut cable, a payment failure — and none of the rungs below will find it.

Rung 2 — read the fault code, including the vendor field

Classify it into one of three: self-clearing, remotely recoverable, or requires a person. If it is the third, stop here and dispatch. Continuing up the ladder on a ground fault wastes time and, if it results in a remote reset, re-energises equipment that tripped for a reason.

Rung 3 — read the message log

What was the unit doing when it failed? A StartTransaction that never received a response looks nothing like a MeterValues stream that stopped mid-session, and they have different causes. This is the rung that most often produces the actual answer, and it is the one most often skipped because it requires the log to be readable by operations.

Rung 4 — check connectivity

Has the unit been sending heartbeats? How often has it reconnected today? A charger that reconnects forty times a day is on a failing link, and every fault it reports is suspect. A silent charger is a network problem, and network problems are fixed differently — you cannot reset a charger you cannot reach.

Rung 5 — change availability

Set inoperative, then operative. This clears more internal state than people expect, interrupts nothing if the unit is idle, and is entirely non-destructive.

Rung 6 — soft reset

Reboots the application without cycling power. Recovers most transient states. Destroys the evidence in rungs two through four, which is why it sits here rather than at the top.

Rung 7 — hard reset

Full power cycle, and the last thing to try remotely — because a unit that does not come back has just converted a probable dispatch into a certain one.

THE HABIT THAT MATTERS MOST

Do not skip to rung six. Resetting first makes the charger work and teaches you nothing; the fault returns next week and nobody knows any more than before. Two minutes on rungs two to four is how a recurring fault becomes a fixed one.

Make the dispatch cheap when it happens

A visit that follows a completed ladder is a different job from one that follows a phone call. The technician leaves knowing the fault code, the message history, what has already been tried and which part is most likely — which usually means they carry the right spare and finish in one visit.

CHAPTER 04

Faults that need special handling

Three fault classes exit the ladder early. Each breaks it for a different reason, and each has caused avoidable harm at operators who treated them as ordinary.

Ground faults and residual-current trips

Protection operated. Something caused leakage current, and the device did exactly what it exists to do. The charger may report it as GroundFailure, or as OtherError with a vendor code, or — on some hardware — barely at all.

This exits at rung two and dispatches. It is never cleared remotely, because clearing it re-energises equipment that tripped for a physical reason you have not yet identified, in a public place, possibly with a vehicle attached.

> A PLATFORM-LEVEL RULE
>
> If your system allows a remote reset on a ground fault, put a confirmation in front of it that names the risk in words. The cost is one extra click. The alternative is a well-meaning operator clearing it three times before anyone visits.

Common causes, in rough order: moisture ingress into the connector or enclosure, a damaged cable, a vehicle-side fault, and a genuine installation problem. The first is seasonal, which is why ground faults cluster after the first heavy rain.

Communication loss

The unit has gone silent. This breaks the ladder because rungs five, six and seven all require the platform to reach the charger, and it cannot.

The diagnostic question is whether the charger is broken or the link is. Two signals separate them: whether other units at the same site are also silent, which points at the link, and the reconnection history before the outage, which points at a degrading connection rather than a sudden failure.

- Whole site silent: the site's connectivity, the router, or the power. Not a charger fault.
- One unit silent, others fine: that unit's modem, antenna, SIM or power supply.
- Rising reconnection count over the preceding days: a marginal link that has finally failed. This is predictable and worth acting on before it becomes an outage.
- Silent immediately after a firmware update: the update, until proven otherwise.

Missing or impossible meter values

The subtlest of the three, because the charger reports itself as perfectly healthy. Energy is being delivered and the readings are absent, frozen, decreasing, or implausible.

It breaks the ladder because it is a billing incident before it is an operations incident, and the ladder is designed for the latter. A session with no defensible meter reading cannot be invoiced honestly.

SYMPTOM	USUAL CAUSE	ACTION
Readings stop mid-session	Meter or comms module	Bill from last good reading, flag the session
Register decreases	Meter reset	Do not bill; investigate the unit
Delta exceeds physical capacity	Unit or scaling error	Cap, flag, do not bill silently
Unit field says Wh but values look like kWh	Firmware bug	Verify against the vehicle, raise with vendor

EVERY ONE OF THESE BELONGS ON A QUEUE

A session that cannot be billed correctly is a defect. Defects that exist only in logs are defects nobody fixes, and they compound into an unexplained gap between energy delivered and revenue collected that gets discovered months later.

CHAPTER 05

Measuring your dispatches

Dispatches are the largest controllable cost in running a charging network, and almost nobody measures them in a form that supports a decision. Recording "connector replaced, one" against a cost centre tells you what you spent. Recording it against a station, a hardware model and a cause tells you what to change.

Tag every visit with four things

01 The station and connector, not just the site.

02 The hardware model and firmware version at the time.

03 The cause, from a closed list rather than free text.

04 The outcome: fixed on site, part replaced, no fault found, or escalated.

A closed cause list is what makes the data usable. Free text produces four hundred unique descriptions of the same six problems.

The single most useful number

Not dispatches per month. The share of faults resolved without a dispatch. It is directly controllable, it improves as your fault dictionary fills out, and it translates straight into money in a way that a raw count does not.

AND THE SECOND MOST USEFUL

No fault found, as a share of dispatches. A high figure means your triage is sending vans at faults that had already cleared, or at serviceability problems — a blocked bay, a cut cable — that a technician cannot fix either.

Reading the distribution

PATTERN	WHAT IT MEANS	WHAT TO CHANGE
Concentrated on one model	A hardware problem	Procurement, not operations
Concentrated on one site	An installation or environment problem	Revisit the commissioning
Concentrated on one firmware version	A regression	Roll back the cohort, tell the vendor
Spread evenly, mostly connectors	Normal wear	Stock the part at the depot
Mostly "no fault found"	A triage problem	The ladder, not the fleet
Rising after a rollout	The rollout	Stop the next cohort

The site-level finding people miss

A connector replacement rate at one site that is triple the fleet average is almost never a hardware defect. It is usually a kerb the cable is dragged over, a bay layout that makes drivers stretch, or a missing holster. All three are cheap to fix once, and expensive to keep replacing.

Reviewing it

Fifteen minutes a month with three questions: which model, site or firmware is over-represented; what share of visits were avoidable; and what one change would have prevented the most of them. Then make that change, and check next month whether it worked.

That loop, run consistently for a year, does more for a network’s economics than any procurement decision available.

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 10 June 2026. This document is general information about running a charging business, not legal, tax or financial advice.