

AVAILABLE



WHITEPAPER

Measuring What Matters

A Practical Methodology for Charging Success

A driver-centric approach to measuring charging success – from customer attempt to successful session.

INSIGHTS FOR A MORE
RELIABLE CHARGING EXPERIENCE

OCTOBER 2026

Table of Contents

1 **ABSTRACT..... 1**

2 **LAY OF THE LAND..... 2**

3 **THE AVAILABLE CHARGING SUCCESS FUNNEL..... 6**

4 **WHAT THE CHARGING SUCCESS FUNNEL REVEALS 18**

5 **FROM MEASUREMENT TO ACTION 22**

6 **REFERENCES I**

7 **APPENDIX II**

Authors

Maren Rehnelt
Stefan Herr
Maximilian Güldenfuß
Leon van Ess
Florian Remke
Flemming Brieger

Published by

Available GmbH
October 2026

1

Abstract

Charging reliability is increasingly being assessed not only by whether charging infrastructure is available, but by whether drivers can actually charge. Yet while the industry is converging on Charging Success as an important outcome, there is still no common methodology for measuring it. Different definitions of an attempt, different starting points and different criteria for success can produce materially different results from the same underlying charging activity.

This paper introduces the **Available Charging Success Funnel**, a driver-centric methodology for measuring Charging Success across the full charging journey. Rather than treating technical messages or charging sessions as proxies for the driver experience, the methodology reconstructs customer attempts and follows their progression through authentication, transaction start, meaningful energy delivery and successful completion. It distinguishes one central outcome metric – the **Driver Charging Success Rate** – from complementary diagnostic metrics that reveal friction before and during charging.

Applied to real-world charging data, the funnel demonstrates why this distinction matters. Failures can occur before a conventional charging session exists; drivers can recover from failed attempts and ultimately charge successfully; and similar headline outcomes can conceal fundamentally different failure patterns. A single success percentage therefore cannot explain the charging experience on its own.

The paper argues that Charging Success should be measured from the driver journey rather than from the data available most conveniently. Making the denominator, starting point and definition of success explicit creates a methodology that is more transparent and comparable – while preserving the granularity operators need to understand **where success is lost and where to investigate next**.

2

Lay of the land

From availability to actual charging success

For years, uptime has been one of the dominant metrics for assessing the reliability of public EV charging infrastructure. It remains an important metric: if a charger is unavailable, a driver normally cannot use it. But availability alone does not tell us whether a driver who arrives at an available charger can actually charge.

Independent field research has demonstrated this distinction. Rempel et al. (2024) tested 655 CCS charging ports across 182 public charging stations in the San Francisco Bay Area and classified a port as functional only if it successfully charged an EV for at least two minutes. Just 73.3% of the tested ports met that criterion, compared with uptime figures of 95–98% reported by charging network operators. The authors concluded that more precise and verifiable definitions of uptime, downtime and excluded time were needed.[1]

The discrepancy is not necessarily a contradiction. **Uptime and Charging Success measure different things.**

Uptime describes whether charging infrastructure is technically operational or available over a period of time. Charging Success asks what happens when someone actually tries to charge.

The distinction is also visible at the regulatory level. The EU Alternative Fuels Infrastructure Regulation (AFIR), for example, requires operators to make dynamic data available on both the operational status of a public charging point and its availability.[2] These are important infrastructure states, but neither establishes whether an individual charging attempt ultimately succeeds.

That outcome depends on an entire chain of interactions. The vehicle and charger need to communicate, authentication and authorization may need to succeed, a transaction needs to start, meaningful energy needs to flow, and charging needs to continue sufficiently to fulfil its purpose.

Research from the U.S. Department of Energy's ChargeX Consortium illustrates how many points of friction can occur along this journey. Its work identified 52 customer pain points, including 18 around starting a charge. These include failed authentication or payment, failure of the EV or EVSE to start charging, repeated attempts, incomplete sessions and lower-than-expected charging power.[3]

Reliability, in other words, is not only a question of whether infrastructure is available. It is also a question of whether charging actually works.

The industry is converging on the outcome - not yet on the methodology

The industry is increasingly responding to this distinction.


ChargeUp Europe has highlighted limitations in the way uptime is measured across Europe and proposed a common methodology, noting that different definitions, exclusions and calculation approaches can make nominally similar uptime figures difficult to compare.[4]

At the same time, several approaches have begun looking beyond infrastructure status towards actual charging outcomes.

Kempower distinguishes uptime from charging success and separates technical charging failures from usability-related issues.[5] Paren's Reliability Index differentiates between successful sessions that work on the first attempt, successful sessions requiring retries, failed charging attempts and station downtime.[6] J.D. Power takes a driver-oriented perspective through "non-charging visits" - instances in which EV owners visit a public charger but are unable to successfully charge.[7]

More recently, P3 Group, together with co-authors from Germany's National Centre for Charging Infrastructure, proposed a dedicated Charging Success Rate as a user-oriented measure of charging reliability.[8] Based on more than 20,000 in-field charging tests across Europe, P3 reports observed charging-session success rates from the user perspective of between 80% and 86% over the preceding 24 months.[9]

Published approaches cover different parts of the charging journey



These approaches differ in their definitions and units of analysis, but they point in the same direction: **reliability is increasingly being evaluated not only through infrastructure availability, but through the outcome of actual charging interactions.**

The relevant question is shifting from:

Was the charger operational?

to:

Was the customer able to charge?

The industry is therefore beginning to converge on the question. **It has not yet converged on a single methodology for answering it.**

Existing approaches illustrate the challenge. A field test may ask whether an EV can charge for a minimum period. A consumer study may ask whether a driver ultimately succeeded during a charging visit. A session-based methodology may evaluate what happened after a charging transaction was created. Another approach may explicitly account for retries before a successful session.

Each perspective can provide valuable information. But they do not necessarily measure the same thing.

Why methodology matters

At first glance, Charging Success Rate appears to be a simple KPI:

Charging Success Rate = Successful charging attempts / Total charging attempts

The complexity lies in the definitions behind that formula.

What constitutes an **attempt**? Does measurement begin with the driver's first interaction, with authentication, or only once a transaction exists?

What constitutes **success**? Is a transaction successful once it starts? Once energy flows? Does a minimum amount of energy or duration need to be reached? And how should a session be classified if charging begins successfully but terminates unexpectedly?

How should **retries** be treated? A driver who succeeds immediately and a driver who succeeds only after several failed interactions may both ultimately charge, but the journey to that outcome is not the same.

These are not minor technical choices. They determine which parts of the charging journey enter the measurement at all — and therefore what the resulting percentage actually represents.

This becomes particularly important when failures occur before a conventional charging session exists. If measurement starts only once a transaction has been successfully created, failed attempts during communication, authentication or authorization can disappear from the denominator. Conversely, a methodology that starts earlier in the journey needs a reliable way to distinguish individual technical signals from actual customer attempts.

The same applies at the other end of the journey. A started transaction does not necessarily mean that meaningful charging occurred, and meaningful charging does not necessarily mean that the session completed without an infrastructure-induced interruption.

A robust Charging Success methodology therefore needs to make explicit decisions about **where measurement begins, what is being counted, and what outcome qualifies as success.**

Without this clarity, Charging Success risks reproducing one of the central weaknesses of uptime measurement: two correctly calculated headline percentages can represent materially different underlying realities.

The objective should therefore not be simply to replace uptime with another headline KPI. It should be to establish a measurement framework that is consistent enough to compare, granular enough to diagnose, and simple enough to communicate.

The industry is increasingly aligned on the outcome that matters:

whether charging works.

The remaining question is how to measure that outcome rigorously.

The following chapter introduces the Available Charging Success Funnel — a methodology designed to do exactly that.

3

The Available Charging Success Funnel

Charging success is simple from a driver's perspective: a driver arrives at a charging station intending to charge, and either succeeds or does not. Measuring that experience from charging infrastructure data, however, is less straightforward.

A single charging journey can generate multiple authentication signals, repeated attempts, transactions that start without meaningful energy delivery, or sessions that begin charging but terminate unexpectedly. Looking at any one of these technical events in isolation can therefore produce a technically correct metric that does not accurately reflect what the driver experienced.


The **Available Charging Success Funnel** addresses this by following the charging journey across its successive stages: from the driver's attempt to authenticate, through the start of a transaction and the delivery of meaningful energy, to the successful completion of the charging session.

The funnel focuses on charging journeys that become observable through charging infrastructure data, primarily OCPP. It therefore does not capture every possible source of driver friction before a charging attempt is recorded — for example, a driver arriving at a charger that was shown as available but is in fact unusable and leaving without initiating an attempt. Such issues can be detected through other data sources and analytical methods, but they sit outside the scope of the Charging Success Funnel itself.

The funnel distinguishes between four complementary measures:

1. **Authentication Success Rate** measures whether a customer attempt results in successful authentication.
- A **Transaction Success Rate** measures whether successful authentication is followed by a charging transaction.
2. **Start Charging Success Rate** measures whether a started session develops into meaningful charging rather than ending prematurely or delivering insufficient energy.
3. **Driver Charging Success Rate** measures whether a started charging session ultimately becomes a successful session without an infrastructure-induced interruption.

The Available Charging Success Funnel



Importantly, these rates are conditional stage metrics. Each rate uses the population that reached the relevant stage as its denominator. They therefore answer different questions and should not be interpreted as interchangeable measures of the same outcome.

Of these measures, the Driver Charging Success Rate is Available's central measure of charging success. It captures the outcome closest to what the driver experiences once a charging session has started: did the session result in meaningful charging and continue without being terminated by an infrastructure-induced interruption? The other metrics preserve the preceding stages of the journey and make visible where friction occurred before or during charging.


Understanding that outcome, however, starts one step earlier: with correctly identifying what constitutes a driver's attempt to charge.

Measuring the driver, not the message

A charging station does not directly record a driver's intention to charge. It records technical events. One driver interaction can generate several such events: repeated card taps, retries after a failed authentication, or a switch from one authentication method to another.

Counting each technical signal as a separate attempt would therefore confuse system activity with driver experience.

From OCPP signals to customer attempts



The Available methodology reconstructs the driver's interaction in three levels:

A **signal** is an individual authentication-related technical message recorded by the charging infrastructure.

Related signals from the same authentication token occurring in immediate succession are consolidated into a **token-try**. This prevents, for example, near-simultaneous duplicate signals or repeated taps from automatically being interpreted as independent driver attempts.

Token-tries are then grouped into a **customer attempt**: one driver's effort to initiate charging at one station. This grouping can span authentication methods. If, for example, a driver plugs in and the initial vehicle-based authentication is unsuccessful, then immediately presents an RFID card to initiate charging, the methodology can treat both interactions as part of the same customer attempt rather than as two unrelated attempts.

The exact grouping logic involves configurable time windows and additional validation rules. These parameters are calibrated to the characteristics of the underlying network and are deliberately not fixed as universal constants: the relevant methodological principle is to reconstruct the customer interaction rather than equate every technical message with a new attempt.

Why timing alone is not enough


Reconstructing customer attempts introduces an important measurement challenge. Two authentication events close together can represent different technical steps within a single driver interaction – for example, an unsuccessful vehicle-based authentication followed by the driver presenting an RFID card – or they can represent two independent driver interactions.

A grouping window that is too narrow risks splitting one continuous driver interaction into several attempts. A window that is too broad risks merging independent interactions. Timing therefore cannot always provide sufficient evidence on its own.

The same time window can represent different customer journeys

One charging station - schematic illustration of the mechanism, spacing not to scale

Failed Accepted | The customer is named here for clarity; only the token and the outcome are in the data.



Where additional information is available, such as the charging connector used, the methodology uses it to improve the attribution of technical events to the underlying customer journey. Where the available data does not allow a distinction to be made reliably, the resulting uncertainty is treated as a limitation of the measurement rather than hidden behind an artificial degree of precision.


This distinction also determines how first-attempt performance is measured. A first-attempt metric evaluates whether the driver's first customer attempt proceeds successfully through the funnel. If a new customer attempt is required before charging ultimately succeeds, the eventual session may still be successful, but the original attempt was not.

The distribution illustrates an important characteristic of real charging behaviour: a failed first interaction does not necessarily mean a failed charging journey. Drivers retry, or switch authentication methods, and can ultimately succeed. Conversely, repeated attempts can be an important indicator of friction even when the final outcome is successful.

This is why the **customer attempt**, rather than the individual technical message, forms the base unit for measuring the first stage of the Charging Success Funnel.

Charging Success is often observed with additional token tries

Bars: share of charging attempts · Line: success rate, Sample dataset across Available network, July 2026



Authentication: did the driver gain access to charging?

The first transition in the Charging Success Funnel is from a customer attempt to successful authentication.

Authentication is considered successful when the customer attempt contains an authentication that is accepted. The **Authentication Success Rate** therefore measures successful authentications relative to customer attempts rather than simply counting accepted and rejected OCPP messages.

$$\text{Authentication Success Rate} = \frac{\text{accepted customer attempts}}{\text{all customer attempts}}$$

This distinction matters. If several technical interactions belong to the same customer attempt, counting them independently can overstate both the number of attempts and the amount of friction experienced by drivers.

At the same time, an aggregated authentication rate alone does not explain why authentication fails.

Available therefore distinguishes between different authentication outcomes and failure reasons. Depending on the available data, denied authentications can, for example, be associated with an invalid or blocked token, an expired authorization, or other identifiable causes.

The composition of these failures can differ across authentication methods. This additional classification turns authentication success from a pure performance metric into a diagnostic measure: the rate indicates **how**

often authentication fails, while the underlying reason structure helps identify **where to investigate why it fails**.

Importantly, absence of evidence is not automatically classified as failure. Where a status is missing or the available technical evidence does not support a clear accepted-or-denied outcome, the methodology does not simply assign the event to the failure population. Similarly, ambiguous outcomes are kept distinct where the underlying data cannot reliably establish a driver-facing failure.


Successful authentication, however, remains only the first gateway in the charging journey. **Being authorized to charge does not mean that charging actually happened.**

From authentication to actual charging

After authentication has succeeded, two further questions determine whether the driver actually begins to charge.

First, did a transaction start? Second, did that transaction develop into a meaningful charging session?

Not every successful authentication leads to effective charging



Both thresholds — minimum duration and minimum delivered energy — are configurable.

These are deliberately treated as separate steps.

A successfully authenticated customer may never see a transaction start. Conversely, a transaction can technically start but end almost immediately or deliver so little energy that describing it as successful charging would misrepresent the driver's experience.

The **Transaction Start Rate** captures the first transition: the share of successful authentications that result in a started session.

$$\text{Transaction Start Rate} = \frac{\text{started sessions}}{\text{successful authentications}}$$

Once a transaction has started, Available qualifies whether it represents meaningful charging. A session must exceed both a defined minimum duration and a defined minimum amount of delivered energy to qualify as an **effective session**.

The two conditions are evaluated separately because they capture different failure patterns. A session can remain active for some time while delivering little or no meaningful energy. Conversely, energy can briefly flow before the session terminates prematurely.

The thresholds used for this qualification are calibrated rather than treated as universal constants. Their purpose is not to define how long or how much every charging session should deliver, but to distinguish meaningful charging from technical transactions that do not represent a meaningful charging experience.

This gives rise to the **Start Charging Success Rate**: the proportion of started sessions that develop into effective charging.

$$\text{Start Charging Success Rate} = \frac{\text{effective sessions}}{\text{started sessions}}$$


At this stage, three outcomes are therefore distinguished: a session may be too short, it may run for sufficient time but deliver insufficient energy, or it may qualify as effective. These categories are applied sequentially along the funnel so that a session is not counted as failing at multiple stages simultaneously.

An effective session has therefore cleared an important hurdle: charging genuinely occurred. But it can still fail from the driver's perspective if the charging infrastructure terminates the session unexpectedly.

Did charging complete successfully?

The final stage of the Charging Success Funnel distinguishes an effective charging session from a successful one.

How session stop reasons are classified



A charging session can end for many reasons. Some endings are deliberate: the driver stops the session locally, remotely, or through an unlock command. Others indicate that charging was terminated

involuntarily, for example by a power loss, forced reset, reboot, emergency stop, or another unambiguous infrastructure-side failure.

These outcomes should not be treated alike. Available therefore classifies whether an effective session was **interrupted**. A deliberate end to charging is not considered an interruption. An unambiguous infrastructure-induced termination is.

Every started session is ultimately assigned to one outcome along the funnel:

1. Too short — the session does not reach the required minimum duration.
2. Insufficient energy — the session runs for sufficient time but does not deliver the required minimum energy.
3. Interrupted — the session is effective, but charging is terminated by an infrastructure-induced interruption.
4. Successful — the session is effective and ends without an infrastructure-induced interruption.

Only the fourth outcome is classified as a successful session. This gives us two important definitions:

- Effective session = a started session that meets the minimum duration and minimum energy criteria.
- Successful session = an effective session that ends without an infrastructure-induced interruption.

The **Driver Charging Success Rate** measures successful sessions relative to started sessions.


$$\text{Driver Charging Success Rate} = \frac{\text{successful sessions}}{\text{started sessions}}$$

This is the central metric in the Available methodology. Rather than asking whether a technical transaction existed, it asks whether a started charging session ultimately provided a meaningful and uninterrupted charging experience.

Detecting interruptions across session boundaries

Real-world charging data can make even this apparently simple question difficult to answer. A technical interruption may end one recorded session and be followed shortly afterwards by another session for the same vehicle. In the underlying data, these appear as separate transactions. From the driver's perspective, however, they can be part of the same interrupted charging experience.

Available therefore uses additional evidence to identify such patterns. Where consecutive sessions can be associated with the same vehicle, occur sufficiently close together and show a plausible continuity in state of charge, the sequence can provide evidence that charging was interrupted, even without respective OCPP signals, and subsequently resumed.



The window, the full-charge threshold and the SoC tolerance are configurable.

The precise thresholds and matching rules form part of the implementation rather than the conceptual framework. The methodological principle is what matters:

Technical session boundaries do not necessarily correspond to boundaries in the driver's charging experience.


The same caution applies to stop reasons. Some reasons provide sufficiently clear evidence of a deliberate stop or infrastructure-induced interruption; others do not. Charger hardware, firmware and protocol versions can produce different technical signals for similar physical events. Where the available evidence does not allow a reliable attribution, the methodology does not manufacture one.

The Charging Success Funnel in practice

Applying these definitions to an exemplary analysis of the Available network demonstrates why the point at which success is measured matters.

The same underlying charging journeys produce different rates depending on the stage being measured:

The Available Charging Success Funnel



These figures do not contradict one another. They answer different questions and use different denominators.

The **Authentication Success Rate** describes whether customer attempts successfully pass authentication. The **Transaction Start Rate** measures the transition from successful authentication to a started transaction. The **Start Charging Success Rate** evaluates whether started sessions become meaningful charging sessions. And the **Driver Charging Success Rate** captures whether started sessions ultimately become successful sessions without an infrastructure-induced interruption.

The **First Try Session Success Rate** provides a complementary end-to-end perspective. It starts with the driver's first authentication attempt and follows that attempt through the entire funnel. Every point of friction — failed authentication, more than one token-try, failure to start a transaction, insufficient duration or energy, or an interruption — can prevent that first attempt from resulting in a successful charging session.

$$\text{First Try Session Success Rate} = \frac{\text{successful sessions resulting from the first customer attempt}}{\text{first customer attempts}}$$

Unlike the conditional stage rates, the First Try Session Success Rate therefore measures the outcome directly from the beginning to the end of the journey. It should not be reconstructed by multiplying the individual stage rates, because the stages use different populations and the links between authentication and session records are not always lossless.

This distinction matters because drivers can recover from failure. A first customer attempt may fail, followed by another attempt that ultimately produces a successful session. The eventual charging experience can therefore be successful even though the journey to get there involved substantial friction.

This is precisely why a single success percentage without a clearly defined denominator can be misleading.

A Charging Success Rate is only meaningful when it is clear what is being counted, where measurement begins, and what outcome qualifies as success.


The Charging Success Funnel makes those choices explicit.

From a KPI to a diagnostic framework

The value of the Charging Success Funnel goes beyond producing a more driver-centric success rate. Because it preserves the individual transitions in the charging journey, it also shows **where success is being lost**.

Charging Success varies across different countries

Sample dataset across Available network, July 2026; Country-level comparison across the analysed charging network



An exemplary breakdown across parts of the Available network illustrates this diagnostic value. The differences are visible well before the final session outcome. In some subsets, a substantial share of customer

attempts is lost during authentication, while in others a much larger proportion passes authentication successfully. Downstream session performance, by contrast, can show a different pattern again.

This matters because similar headline success rates do not necessarily point to the same underlying problem. A loss at authentication requires a different investigation from a transaction that fails to start, a session that delivers no meaningful energy, or a charging session that terminates unexpectedly.

The funnel therefore provides a starting point for diagnosis: it identifies at which stage of the charging journey success is lost. Determining why that loss occurs requires the next layer of analysis, using the underlying technical signals and additional available data.

The purpose of the country breakdown shown here is not to establish a market benchmark or rank countries. It illustrates how an aggregate charging success rate can conceal fundamentally different failure patterns underneath.

For an operator, that distinction is actionable. Improving authentication requires a different investigation from improving transaction starts, reducing ineffective sessions, or preventing interruptions.

The Charging Success Funnel therefore changes the question from:

What is our charging success rate?

to two more useful questions:

Did the driver successfully charge?

and

If not, where in the journey did success
break down?

Measuring the entire driver journey makes both questions answerable — and makes Charging Success a metric that can be used not only for reporting performance, but for improving it.

4

What the Charging Success Funnel Reveals

The Charging Success Funnel does more than produce a more granular set of metrics. Applied to real charging activity, it changes what becomes visible.

Looking only at the outcome can hide friction along the way. Looking only at started sessions leaves out attempts that never became sessions. And looking at one aggregate rate can conceal the stage at which success was lost. The value of the funnel lies in revealing **how success is achieved – and where it is lost**.

The findings below draw on Available's analysis of charging activity across different customers in July 2026. They describe this sample, not the charging market as a whole. A separate analysis of token attempts covers 16 customers and is identified where used.

A substantial share of observed loss occurs before a session starts

In the pooled sample, **20.9% of recorded authentication attempts did not result in successful authentication**. Once authentication was successful, **96.9% led to a started session**.

This places a substantial share of the visible loss before a session exists. A metric based only on started sessions can describe what happened after that point. It cannot show the attempts that failed before it.

A failed authentication attempt is not necessarily a driver who ultimately failed to charge. A driver may try again or change authentication method, and several technical attempts may belong to one charging journey. But those attempts are still part of the experience. Even when the driver eventually succeeds, a session-based metric alone does not show the difficulty of getting started.

The first finding is therefore not simply that authentication matters. It is that **the starting point of measurement determines which difficulties can be seen at all**.

This does not mean that optimizing authentication is automatically the best lever for any CPO. Many failed attempts are outside the CPO's control, for example when a customer is simply not authorized to charge at a given station. In addition, OCPP itself limits how precisely failure causes can be distinguished: a rejected authorization does not always reveal whether the reason was a technical fault or a backend issue.

A successful charge can still be a poor driver experience

Consider a driver whose first attempt to authenticate via RFID fails. They try another method, successfully start a session and complete the charge. Viewed from the final session, the outcome is successful. Viewed from the beginning of the journey, it required the driver to recover from an initial failure.

The token-attempt analysis shows that such friction is more than a hypothetical possibility. Across the analysed attempts, there were **1.55 token tries per attempt** on average, and **23.5% of attempts involved more than one authentication token**. Only **54.5%** followed the simplest observed pattern: one try with one method.

Attempts involving additional effort could still end successfully. The success rates within the chart's buckets were **89.6%** for "one try plus a second method" and **91.4%** for "two tries" with the primary method. These figures do not mean that an additional try or a change of method causes success. They show that the final outcome can coexist with a more difficult path to reaching it.

Recovery masks friction. A driver may ultimately charge while having had to repeat an action, change method or work around a failed first attempt. A network could improve that experience without necessarily changing its rate of ultimately successful sessions.

This is why **First Try Session Success** and **Driver Charging Success** offer different perspectives. The former asks whether the journey succeeded from the first customer attempt. The latter measures successful sessions relative to started sessions. Their percentages have different denominators and must not be treated as a numerical measure of how many drivers recovered. Together, however, they make clear why the final session outcome cannot tell the whole story.

Losses after session start are not evenly distributed

The later stages of the pooled funnel reveal another distinction: starting a session does not yet mean that effective charging has taken place.

Started sessions represent **76.6%** of the funnel's initial authentication-attempt baseline. The funnel then deducts **5.2 percentage points** for short sessions and **2.2 percentage points** for low-energy sessions. Short and low-energy sessions are defined as mutually exclusive categories in the funnel. This leaves **69.3% effective sessions**. A further **0.6 percentage points** are deducted for interrupted sessions, leaving **68.7% successful sessions**. The deduction of the interrupted sessions is a lower bound: with the detection rules applied here, interruptions that are not clearly identifiable remain within the successful sessions, so the true share of interrupted sessions is higher.

The deduction for short sessions is more than four times as large as the deduction for low-energy sessions. Loss through interruption after effective charging is smaller still. In this sample, the larger visible loss after session start is associated with reaching effective charging, rather than with the interruption of a session that has already become effective.

These figures do not establish why sessions were short. Nor do the successive deductions reveal how the underlying short- and low-energy categories would overlap if examined independently. What they show is that **short duration, low energy delivery and interruption contribute very differently to the measured result**. Combining them into one category of “failed sessions” would hide that pattern.

Similar outcomes can conceal different failure profiles

A headline outcome describes how many journeys reached a defined endpoint. It does not, by itself, explain where the others were lost.

Two network subsets could reach a similar final success rate through very different paths. In one, drivers might struggle mainly before a session starts. In another, sessions might start reliably but fail to become effective charging. Those patterns would call for different investigations even if the headline percentages looked alike.

The country analysis illustrates why the stages matter. Its first-try panel shows wide differences in how far attempts progress through the early journey. Among countries with meaningful attempt volumes, the share qualifying as first-try ranges approximately from **20% in Ireland** to **90% in Austria**. In the separate session-outcome panel, results for starting charging and delivering energy are much closer together across most countries.

The panels have different denominators: **attempts** in the first-try view and **sessions** in the session-outcome view. They cannot be read as a direct comparison of the same success rate, and the chart does not establish that two countries have identical Driver Charging Success Rates. It does demonstrate how a view focused on sessions can show much less variation than a view that also includes the earlier attempt journey.

The broader diagnostic point holds: **a similar outcome does not imply a similar path to that outcome**. To understand what needs attention, operators need to see the losses at each stage rather than compare the final number alone.

The denominator can change the story

The pooled analysis produces several rates from the same charging activity. **Authentication Success** is **79.1%** of authentication attempts. **Transaction Start** is **96.9%** of successful authentications. **Driver Charging**

Success is 89.7% of started sessions. The chart also reports a First Try Session Success Rate of 30.6%, applying a first-try perspective earlier in the journey.

These percentages differ because they answer different questions. None becomes meaningful without knowing its unit of measurement, where its denominator begins and what qualifies as success.

A rate based on started sessions cannot describe all attempts to begin charging. A rate based on authentication attempts cannot, on its own, describe whether meaningful energy was delivered. A rate based on effective sessions says something different again about uninterrupted completion.

The problem is not that one of these rates is necessarily wrong. The problem is calling them all “Charging Success Rate” without specifying what they measure.

This matters when comparing results across networks, countries or published methodologies. Two percentages can be calculated correctly and still be answers to different questions. The funnel makes those questions explicit before the figures are interpreted.

Better measurement changes what you can see

A single outcome measure remains useful. Operators need a clear answer to whether charging succeeded for the journeys included in that measure. The Driver Charging Success Rate provides that central outcome for started sessions in Available’s methodology.

But the findings in this chapter show what the number alone cannot explain. Substantial visible loss occurs before session start. Some successful outcomes require repeated effort. Among started sessions, the observed losses are concentrated in particular stages. And different definitions can produce very different percentages from the same activity.

The funnel retains the clarity of an outcome measure while showing the stages behind it. That does not mean every stage must become an equally prominent KPI, or that the funnel itself establishes the root cause of every failure. It gives the outcome a diagnostic context.

For the driver, the essential question remains simple: **they came to charge. Did it work?** The difficulty lies in reconstructing the full journey from technical data. By connecting that journey to the final outcome, the funnel turns Charging Success from a number that describes performance into a framework that helps explain it.

Once operators can see **where** success is lost, the next question becomes practical: **what should they investigate and improve?**

5

From Measurement to Action

A more precise way of measuring Charging Success is not an end in itself. Its value lies in what it enables operators to see — and, ultimately, what it allows them to improve.

A single success percentage can indicate that a problem exists. But it rarely explains where that problem occurs or what kind of intervention is required. By breaking the charging journey into distinct stages, the Charging Success Funnel turns Charging Success from a reporting metric into a diagnostic framework.

For CPOs, this has four practical implications.

Start with the driver journey, not the available data

Charging Success is often defined from the data backwards: operators look at the events and fields available in their backend systems and construct a KPI from what can be observed most easily.

The Charging Success Funnel takes the opposite approach. It starts with the experience that needs to be understood.

A driver arrives at a charge point with a simple intention: to charge. From that perspective, the relevant journey begins before a charging session has necessarily been created in the backend. Authentication attempts that fail, repeated attempts to initiate a charge, or a switch from one authentication method to another are part of the driver experience even if they never result in a conventional session record.

This distinction matters. When measurement begins only with successfully started sessions, everything that went wrong before that point disappears from the denominator. The resulting metric may accurately describe the performance of sessions that were created, while saying considerably less about the driver's ability to start charging in the first place.

The first question should therefore not be [What data do we have?](#) but:

What part of the driver journey are we trying to measure — and which technical signals allow us to reconstruct it?

This shift in perspective is fundamental. Technical data remains the basis of measurement, but the driver journey determines how that data should be interpreted.

Define the denominator before discussing the rate

A percentage becomes meaningful only when it is clear what sits behind its numerator and denominator.

As the Charging Success Funnel demonstrates, several success rates can be calculated from the same underlying charging journeys. Each may be mathematically correct while answering a different question. A rate based on started sessions measures something fundamentally different from a rate that includes customer attempts that never resulted in a session.

For that reason, comparisons of Charging Success should begin with the definition rather than the percentage itself.

Three questions for any Charging Success Rate

1

What is the unit of measurement?

Sessions, transactions, authentication attempts, customer attempts, charge points - or something else?

2

Where does measurement begin?

At the driver's first attempt to charge, after successful authentication, or only once a session has started?

3

What counts as success?

A started transaction, energy delivery, a session above a defined minimum threshold, or successful completion without interruption?

Without answers to these questions, two percentages labelled "Charging Success Rate" may not be comparable at all.

Greater transparency around these definitions would therefore do more for meaningful industry comparison than simply reporting more Charging Success percentages.

Measure both outcome and friction

Ultimately, the most important question is whether the driver was able to charge. This is why the **Driver Charging Success Rate** serves as the central outcome metric within the Charging Success Funnel.

But outcome alone does not capture the entire experience.

A driver may ultimately charge successfully after a failed authentication attempt, a retry, or a change in authentication method. From an outcome perspective, the journey was successful. From an experience perspective, however, it contained friction.

This is why outcome and friction should be measured separately rather than collapsed into a single metric.

The Driver Charging Success Rate captures whether a started charging journey ultimately resulted in successful charging. Complementary indicators — such as First Try Session Success, Authentication Success and the number of attempts required — reveal how easily the driver reached that outcome.

This distinction also avoids turning Charging Success into an ever-growing dashboard of equally weighted KPIs. There should be a clear hierarchy: **one central outcome metric**, supported by diagnostic metrics that explain the friction behind it.

The objective is not to maximize the number of success rates being reported. It is to understand both whether charging worked and how difficult it was for the driver to make it work.

Use the funnel to react

Once Charging Success is decomposed into stages, the operational question changes.

Instead of asking [Why is our Charging Success Rate low?](#), an operator can ask [Where in the charging journey are we losing successful outcomes?](#)

Losses before or during authentication point toward a different set of potential causes than losses between authentication and transaction start. Sessions that start but fail to become effective require investigation of yet another part of the charging process. Interruptions during otherwise effective charging sessions indicate a different reliability problem again.

The funnel therefore creates a direct connection between measurement and investigation:

Where success is lost determines where to look

QUESTION	MEASURED BY	IF LOW, LOOK AT
Did the attempt pass authentication?	Authentication Success Rate	Authentication flows
Did a transaction start?	Transaction Start Rate	Session initiation
Did the session deliver meaningful energy?	Start Charging Success Rate	Energy delivery
Did charging complete without interruption?	Driver Charging Success Rate	Reliability during charging

The funnel itself does not determine the root cause. A failed authentication, for example, can have multiple technical or operational explanations. Nor does every unsuccessful journey necessarily indicate a charge-point fault.

What the funnel does is narrow the search space.

It identifies the stage at which the observed driver journey diverged from the intended one. Operators can then combine this information with technical diagnostics, error codes, charge-point telemetry and operational data to investigate the underlying cause.

This is the practical value of more granular Charging Success measurement: **different failure stages require different operational responses.**

A single percentage can tell an operator that performance changed. A driver-centric funnel can help show where to look next.

Ultimately, better measurement does not improve Charging Success by itself. But it changes the quality of the questions an operator can ask:

from

How successful is my network?

to

Where are drivers struggling, and what
should we investigate to improve their
experience?

6

References

- [1] Rempel, D., Cullen, C., Bryan, M. M. & Cezar, G. V. (2024). *Reliability of Open Public Electric Vehicle Direct Current Fast Chargers*. Human Factors, 66(11), 2528–2538. DOI: 10.1177/00187208231215242.
- [2] European Union (2023). *Regulation (EU) 2023/1804 on the deployment of alternative fuels infrastructure (AFIR)*, Article 20.
- [3] National Renewable Energy Laboratory (2024). *Impact of Electric Vehicle Charging*. NREL/TP-5400-89896; drawing on the ChargeX Consortium customer pain-point framework.
- [4] ChargeUp Europe (2024). *Understanding uptime in the EV charging industry: towards a common methodology*.
- [5] Kempower (2024). *Charging Success Rate and Charger Uptime in Public EV Charging*; accompanying published charging-performance data based on charging sessions recorded through ChargeEye.
- [6] Paren (2025–2026). *U.S. EV Fast Charging State of the Industry* and Reliability Index methodology.
- [7] J.D. Power (2025). *U.S. Electric Vehicle Experience (EVX) Public Charging Study*.
- [8] P3 Group & co-authors from the National Centre for Charging Infrastructure (2026). *Charging Success Rate*. See also National Centre for Charging Infrastructure, “*Charging Success Rate misst Erfolg von Ladevorgängen nutzerorientiert*”, 24 March 2026.
- [9] P3 Group (2026). *Charging Success Rate*. P3 reports analysis of more than 20,000 in-field charging tests in Europe and an observed charging-session success rate from the user perspective of 80–86% over the preceding 24 months. P3 cites its *Power2Drive Europe Conference 2024* work, *Importance of Interoperability for Reliability of Public High-Power Charging*, as the source for this analysis.

7

Appendix


Limitations

The results in this paper are subject to five limitations:

The funnel only sees what reaches OCPP

Measurement begins with the first authentication signal a charging station transmits via OCPP. Failures before that point — a charger shown as available but unusable, a defective card reader, a driver leaving without interaction or worse, a stolen cable— produce no observable attempt. **The reported rates therefore overestimate Charging Success, with respect to these failures, they are an upper bound.** OCPP data alone cannot quantify this gap.

What the funnel can see




Attempt reconstruction is a heuristic

Customer attempts are reconstructed using time windows and validation rules (Appendix). No ground truth confirms which token-tries belong to the same driver. **The resulting rates depend on these parameters:** a wider attempt window, for example, merges failed and subsequent successful token-tries and raises the

Authentication Success Rate. Rates based on customer attempts are only interpretable together with their parameters.

Rates depend on the attempt window




One parameter set across all networks

The methodology calibrates its parameters per network, because hardware, firmware, authentication mix and driver behaviour differ. This paper applies a single parameter set to all networks. It fits some use cases better than others, and the rates of other cases are biased accordingly. Differences between countries and networks in Chapter 4 therefore partly reflect parameter fit, not only charging performance. Per-network calibration would reduce this bias, but rates across networks would no longer rest on identical definitions.

One duration threshold cuts each network differently

Session duration of non-interrupted sessions, one line per network · July 2026



The sample does not benchmark the market

The findings are based on one month (July 2026) of data from different Available customers (1 6). The sample is not random and does not capture seasonal effects. Additionally pooled results are dominated by high-volume networks, and some country results rest on low volumes. The results illustrate the methodology, they do not benchmark the market, individual countries or operators.

Published figures are not directly comparable

Published figures such as the 80–86% reported by P3 [9] are based on in-field tests with a different unit of analysis, starting point and success criterion. They cannot be compared directly with the rates in this paper.

These limitations concern the absolute level of the reported rates more than the stage structure of the funnel. **Every reported rate should be accompanied by its data basis and parameters.**

Parameter Values Used in This Paper

Every parameter of the methodology is configurable and is calibrated per network rather than fixed by the methodology itself. The values below are the ones applied throughout this paper.

Parameter	Value used
Retry window (token-try collapsing)	3 seconds
Attempt window (customer attempt grouping)	120 seconds
Minimum session duration	120 seconds
Minimum delivered energy	3 kWh
Interruption linking window	2 hours
Full-charge threshold	80 % SoC
SoC tolerance	1 %