GUIDE

MEASURING A CALL TEAM FROM THE LEDGER

Speed to lead, connection rate, call duration, schedule rate, and show rate: how each is defined, where the number comes from, and how to coach with it.

AUGUST 7, 2026 · 8 MIN READ

Most call-center metrics are reported. An agent logs a call, a manager tallies a sheet, a CRM field gets updated when someone remembers. The number in the Monday meeting is a claim about what happened.

DialSplice measures differently. Every call, callback, and booking is a row in the ledger, so the numbers come from what happened, not what was reported. Speed to lead, connection rate, call duration, schedule rate, and show rate are tracked per agent, per location, and per queue view. Managers coach from the same figures agents see.

This guide defines each of the five metrics, says which ledger events produce it, and describes how to coach with it. It is written for the call-center director who owns the numbers and the manager who has to use them on Tuesday.

Why the ledger

The audit ledger is hash-chained and append-only. Every event links to the previous by SHA-256, so changed, missing, or reordered history is detectable. It stores events and hashes, not message bodies or recordings.

That design was built for compliance. It happens to be exactly what a measurement system needs. A metric read from the ledger has three properties a reported metric does not:

- It cannot be back-filled. A callback claimed at 4:07 PM was claimed at 4:07 PM.
- It cannot be double-counted. One canonical work item has one set of events, whatever view it was worked from.
- It is the same number for everyone. The leaderboard uses the same rows as the compliance export.

The analytics surface reads these rows and presents them by agent, by location, and by view. Nothing in it is entered by hand.

Speed to lead

Definition

Time from a new lead to the first attempt. Measured per work item, reported as a median.

Where the number comes from

A new lead arrives as a routed event: a web form, an inbound message, or a missed-call webhook from your CRM or telephony provider. That event creates a work item with a timestamp and a location. The first attempt is the first click-to-call intent, SMS send, or email send on that work item, each of which is a ledger row with its own timestamp. Speed to lead is the difference.

The clock starts when the event is routed, not when an agent claims the item. Time spent sitting in the Power Hour view counts. That is the point.

How to coach with it

- Read it by view first. Slow speed to lead across Power Hour is a staffing or priority problem. Slow speed to lead for one agent is a coaching conversation.
- Read the median, not the mean. One lead that sat overnight will distort an average and hide a team that is otherwise fast.
- Pair it with the preflight. If an attempt was blocked by quiet hours or a channel block, the delay was policy, not the agent. The ledger records the preflight evaluation, so you can tell.
- Set a target per location, in local time. A lead that arrives at 5:50 PM cannot be worked in four minutes if the location closes at 6:00 PM.

Connection rate

Definition

Attempts that connected, as a share of all attempts.

Where the number comes from

Every attempt is a click-to-call intent in the ledger. Every outcome is a structured disposition on the same work item. An attempt connected if its disposition is one that requires a live contact: “Callback requested,” “Booked,” or “Do not contact.” It did not connect if the disposition is “No answer,” “Voicemail,” or “Wrong number.” The rate is connected attempts divided by all attempts.

The disposition list is yours, configured per location. Which dispositions count as connected is part of that configuration and is versioned with it.

How to coach with it

- Do not coach on the raw rate alone. A low connection rate on a fifth attempt is expected; the cadence is already spacing those out. Read connection rate by attempt number when you can.

- Look at “Wrong number” separately. A rising wrong-number count is a data-quality problem in the CRM, not an agent problem.
- Compare by location and time of day. If one location connects at half the rate of another with the same agents, look at its quiet hours and caller ID, not its people.
- Use it to set the attempt cap. When connection rate on later attempts falls near zero, the cap is doing its job.

Call duration

Definition

Connected-call time, per agent and per outcome.

Where the number comes from

DialSplice does not carry the call. Your telephony provider does. What DialSplice reads from the provider is the call event: ringing, answered, missed, and duration. No audio, no recording, no transcript, and no parties or numbers from metrics endpoints. The duration on the call event is joined to the work item and its disposition in the ledger.

Because the disposition is on the same work item, duration can be reported per outcome: how long a “Booked” call takes compared with a “Callback requested” call.

How to coach with it

- Duration is evidence, not a score. A long call that booked is different from a long call that ended in “No answer.”
- Read it per outcome before per agent. If one agent’s booked calls are twice as long as the team’s, that is a script conversation. If their no-answer calls are long, that is a workflow conversation.
- Pair it with versioned guidance. Scripts are versioned and approved before agents see them. When guidance changes, compare duration by outcome before and after the revision.

- Do not use it as a proxy for quality. Quality scoring needs the audio, and the audio stays in your telephony provider or your QA tool. DialSplice does not do it.

Schedule rate

Definition

Connections that booked, as a share of connected attempts.

Where the number comes from

A booking is a write to your system of record, made from the workspace with identity resolved first. It is journaled as an external effect, once, with the CRM updated from the same event. The disposition “Booked” lands on the work item. Schedule rate is booked work items divided by connected work items.

Because the booking is a journaled effect and not a self-reported field, a “Booked” disposition with no booking event is visible as a discrepancy, and a booking in reconciliation is counted when it resolves, not before.

How to coach with it

- This is the metric that most directly reflects the conversation. Coach it with the script beside it.
- Read it per location and per view. Inbound callers who already asked for an appointment will book at a higher rate than Power Hour leads. Comparing the two is not useful; tracking each over time is.
- Watch the ratio between “Callback requested” and “Booked.” An agent who sets many callbacks and books few may be deferring the ask. The callback SLA playbook covers what happens next.
- Tie it to the downstream funnel. Schedule rate is the middle of Attempts → Connected → Scheduled → Showed. A good schedule rate with a poor show rate means the booking is happening but the commitment is not.

Show rate

Definition

Appointments kept, as a share of appointments scheduled.

Where the number comes from

Your system of record owns the appointment. DialSplice reads appointment outcomes back from it: kept, cancelled, and no-show. Each outcome is tied to the booking event that created the appointment, which is tied to the call that produced it. Show rate is kept appointments divided by scheduled appointments.

This is the only one of the five metrics that comes from outside the call. It is read back from your EHR, not entered by an agent, and it closes the loop between the call team and the clinic.

How to coach with it

- Show rate is the measure of whether the call team’s work mattered. Put it next to the calls that produced it, per agent and per location.
- A low show rate with a high schedule rate usually means appointments are being set too far out, or confirmation steps are being blocked by policy. Check the workflow: a confirmation SMS runs only if the preflight passes at send time.
- Compare by location before by agent. A location with a long wait for the next available slot will have a lower show rate whatever the agents do.
- Use it to settle the argument about outbound. Appointments kept from outbound calls, read from the EHR, is the number a clinic director will believe.

Reading all five together

The analytics surface shows the five metrics as a row of tiles, a bar chart of speed to lead by agent, and a funnel labeled “Downstream impact.” The funnel is the five metrics in order: Attempts, Connected, Scheduled, Showed. In the product example, 412 attempts produced 156 connections, 42 bookings, and 34 kept appointments.

Read the funnel top to bottom and ask one question at each step:

1. Attempts: is speed to lead inside the target for each location?
2. Connected: is connection rate stable by attempt number, and is wrong-number low?
3. Scheduled: is schedule rate moving with the script revision?
4. Showed: are kept appointments rising with bookings, or diverging?

Filter by location first, then by view, then by agent. A problem that appears at the location level is policy or staffing. A problem that appears only at the agent level is coaching.

What the ledger cannot tell you

- What was said. There are no recordings and no transcripts. Quality scoring stays in your telephony provider or QA tool.
- How agents feel about a script. That comes from the non-PHI feedback triage, not from the numbers.
- Whether the lead was any good. The ledger records what happened to it, not where it came from beyond the routed event.

Taking the numbers with you

The same rows behind these five metrics can be exported to your data platform on AWS, Microsoft, or Google Cloud as content-free facts: call aggregates per location and view, work-state transitions with timestamps, appointment outcomes, and analytics facts. Every export batch is a ledger event with a hash. Contact identity, message bodies, and recordings never leave.

For a weekly review without a warehouse, the call log exports to CSV without message content: time, direction, location, agent, outcome, and sync state.

Whichever route you take, the number you present is the number the ledger holds. That is the difference between measuring the work and reporting on it.