



THE CALLBACK SLA PLAYBOOK

How to define a callback window per location, who owns a deferral, what an attempt cap should be, and what to do with the ones that slip.

AUGUST 21, 2026 • 7 MIN READ

A callback is a promise. A lead asked to be called at 4 p.m. local. In most call centers that promise lives in a note, in a CRM, in someone's head. Nobody owns it, so nobody makes it.

In DialSplice a callback is a work item. It is owned by one agent and scheduled in the location's local time. It stays ineligible until its not-before time, then surfaces in the Callbacks view. If it drifts toward its SLA, the deferral is routed to a manager for approval instead of quietly slipping. Retries follow a set cadence with an attempt cap.

This playbook is for the call-center director and the operations owner who have to choose those numbers. It covers what to decide, in what order, and what to do when a callback is missed anyway.

Terms

- Callback window: the time between when a callback becomes eligible and when it is late. Set per location.
- Not-before time: the earliest a callback may surface. A callback is ineligible until then.
- Retry cadence: the spacing between attempts after a no-answer, for example 2h → 4h → next business day.

- Attempt cap: the maximum number of attempts on one work item before it is held for a person.
 - Deferral: moving a callback past its window. Needs a manager's approval.
 - Hold: a work item that has hit its cap or returned an ambiguous provider outcome. It waits for a person, never a retry.
-

Procedure 1: Define the callback window per location

The window is the SLA. It has to be short enough that the contact is still expecting the call and long enough that a location with two agents can meet it.

1. List every location with its time zone and business hours. Use the readiness worksheet if you have filled it in.
2. For each location, pick the latest local time at which a same-day callback may still be attempted. This is the end of the window. It is usually the location's close, or an hour before it.
3. Decide how much slack a callback gets around its scheduled time. A callback scheduled for 4:00 PM local with a one-hour window is late at 5:00 PM.
4. Write the rule as one sentence. For example: deferral past 5:00 PM is routed to a manager.
5. Check the window against quiet hours. A window that ends inside quiet hours cannot be met; the preflight will block the call.

Do not set one window for the whole organization. Locations in different time zones with different hours need different windows, and the configuration is per location anyway.

- ☐ Every location has an end of window in local time.
- ☐ No window ends inside that location's quiet hours.
- ☐ The rule is written as a sentence a manager can repeat.

Procedure 2: Set the retry cadence

The cadence decides what happens after a no-answer. The preflight re-checks consent, quiet hours, and the attempt cap at each step, so the cadence only needs to describe spacing.

1. Choose the first retry interval. Two hours is common; it lets a contact finish whatever they were doing.
2. Choose the second. Widen it. Four hours after the first retry reaches a different part of the day.
3. After that, move to the next business day. Calling the same contact three times in one afternoon is what the cap is there to stop.
4. Write the cadence as a chain: 2h → 4h → next business day.
5. Confirm the chain respects holidays. Next business day means the location's next business day, not the calendar's.

The cadence is part of the location configuration. It is versioned, approved, and can be rolled back. If a location needs a different cadence, give it one; do not bend the rule for everyone.

Procedure 3: Set the attempt cap

The cap is the number of attempts before a work item stops retrying and is held for a person.

1. Count the attempts your cadence produces. A chain of 2h → 4h → next business day, run through twice, is five attempts including the first.
2. Set the cap at or below that count. Five is the reference value in this playbook. Fewer is fine. More needs a reason your compliance lead will sign.
3. Decide what a held item looks like on the manager desk. Held items appear under "Attempt caps reached." A manager releases or closes them.

4. Decide whether a held item can be reopened, and by whom. The answer is usually: a manager, once, with a disposition.

The cap is per work item. A contact with two open work items is a data problem, not a cap problem; the repeat-lead rule updates the open opportunity rather than creating a second one.

- ☐ Cadence written as a chain.
- ☐ Cap set, and not higher than the cadence produces.
- ☐ A manager knows where held items appear.

Procedure 4: Assign ownership

One agent owns a callback. That agent's name is on the work item, and the callback surfaces in that agent's Callbacks view at its not-before time.

1. The agent who sets the disposition "Callback requested" owns the callback by default. That is the person the contact spoke to.
2. If the owner is off shift when the callback is due, the manager moves it. This is a Coverage assignment; the work item is reassigned, not copied.
3. Unworked items at close become an overnight batch with fairness evidence attached. A manager reviews and releases the batch in the morning; nothing is assigned without that review.
4. Two agents never hold the same work item. A single-winner lease is released or expires; there is no shared ownership.

Write down who moves callbacks when an owner is out. If the answer is "whoever notices," the callback will be missed.

Procedure 5: Handle a deferral

A deferral is what happens when a callback cannot be made inside its window. It is the step most teams get wrong, because the easy thing is to let it slide.

1. The agent opens the work item and chooses "Defer → manager." "Keep schedule" and "Release" are the other options. There is no snooze.
2. The manager sees the deferral under "Deferrals awaiting approval" with the original scheduled time and the attempt count.
3. The manager approves or declines. Approval moves the not-before time; the callback stays owned by the same agent unless the manager reassigns it.
4. The approval is logged. The ledger records who deferred, who approved, and when, as events and hashes.
5. A declined deferral goes back to the owner with the original time. The agent makes the attempt or releases the item.

Set a rule for how many times one callback may be deferred. Once is defensible. Twice is a pattern. A third deferral should close the item with a disposition, not push it again.

- ☐ Managers know where deferrals appear and that approvals are logged.
- ☐ A maximum number of deferrals per callback is written down.
- ☐ Agents know that "Defer → manager" is the only way to move a callback past its window.

Procedure 6: Handle the ones that slip

Some callbacks will be missed. The point is that a missed callback is visible, owned, and closed on purpose.

1. Overdue callbacks show in the Callbacks view and count as SLA risk in queue health. "2 at SLA risk" on the manager desk means two work items, not a percentage.

2. The manager decides, for each: attempt now, defer with approval, or close with a disposition.
 3. A callback closed without contact gets "No answer" or "Voicemail," not "Callback requested" again. Re-dispositioning as a callback restarts the cadence and hides the miss.
 4. If the provider returned an ambiguous outcome (the call intent was sent but no event came back), the item is in reconciliation. It is not retried. A manager releases it once the outcome is known.
 5. At the end of the day, review what slipped by location. The call log lists each call with its outcome and sync state, and exports to CSV without message content.
-

Procedure 7: Review weekly from the ledger

Every callback claim, deferral, approval, and attempt is a ledger row. Review from the rows, not from memory.

1. Pull callbacks kept versus scheduled, per location and per agent.
 2. Pull deferrals requested and approved, per manager. A manager who approves every deferral is not reviewing them.
 3. Pull attempt caps reached. A rising count usually means the cadence is too tight or the leads are wrong, not that agents are slow.
 4. Adjust the configuration as a new revision. It is approved, activated, and can be rolled back. Do not adjust it in a meeting and hope agents remember.
-

Reference configuration

The values below are the ones used in the product examples. They are a starting point, not a recommendation for your locations.

SETTING	VALUE
Retry cadence	2h → 4h → next business day
Attempt cap	5
Deferral rule	Deferral past 5:00 PM is routed to a manager
Not-before	The scheduled local time; ineligible until then

What this playbook does not cover

DialSplice does not dial. It does not ring the phone; your telephony provider does. It decides which callback is eligible, who owns it, and what happens when it is missed, and it records all of that. Predictive or power dialing, if you use it, stays with your provider and outside this playbook.