

Reference Material - Survox Dialer

How the Survox Predictive Dialer Works

Overview

The Survox dialer is not a simple "ratio dialer" where you set a fixed number like "dial 3 calls per agent." It is a predictive dialer driven by a live, dynamic algorithm that constantly adjusts how aggressively it dials based on what is happening on the floor in real time. Understanding how it works helps explain why contact rates change throughout a shift, why adding more agents does not always mean more dials, and why some nights perform better than others even with the same numbers.

The Algorithm Inputs

The dialer does not run on a fixed timer. Instead, it recalculates every time a job determines that more numbers need to be dialed -- typically because agents are waiting and available. Each time that happens, the following inputs are evaluated for that job:

1. Drop Rate (Actual vs. Target) -- The dialer tracks a configured drop threshold expressed as a value per 10,000 connects. A setting of 350, for example, means the system will tolerate up to 3.5% of connects resulting in a drop. The dialer does not simply try to stay under this number -- it actively tries to dial as aggressively as possible while landing just below the drop target. The live drop rate updates every time the algorithm fires.

When the drop rate exceeds the configured target, the dialer does not make a small adjustment -- it throttles back aggressively, dropping to near 1:1 dialing until the drop rate falls back below the threshold. Once recovered, it begins climbing back toward the 3:1 cap again. This recovery behavior can cause noticeable slowdowns on the floor and is why sustained high drop rates are worth monitoring closely.

2. Number of Available Agents -- The dialer evaluates how many agents are waiting and not currently on a call for that job. This is what triggers the recalculation in the first place. No

waiting agents means no new dials.

3. Real-Time Connect Rate -- The connect rate also updates every time the algorithm fires. When the contact rate is high, agents are busy more often, the dialer does not need to push as hard to keep them occupied, and drop risk is lower. When the contact rate is low, agents are idle more often, the dialer pushes closer to the 3:1 cap to compensate, and the chance of a drop increases.
4. Capacity Settings (Two Hard Limits) -- Regardless of what the algorithm calculates, two hard caps always apply together:
 - Max 3 calls per waiting agent
 - Max 30 calls per second system-wide

Each job runs this calculation independently for itself. The shared 375-channel pool only becomes a factor if the combined math across all running jobs -- total active agents across all jobs multiplied by 3 -- would push past the 375-channel ceiling. In most normal operating conditions, that ceiling is never reached, and each job dials freely without any competition from other jobs.

What a "Drop" Actually Means

This is one of the most commonly misunderstood parts of how the dialer works.

A dropped call is NOT a respondent who picks up and then hangs up before reaching an agent. A dropped call is the dialer itself hanging up early -- before the call is ever answered -- because the algorithm has determined there is a risk that no agent will be available if the respondent picks up.

Even if the project is configured for 4 rings before abandoning, the dialer may cut a call off at 1, 2, or 3 rings if it calculates that sending that call through to a live pickup would create an agent availability problem. The respondent's phone may ring once or twice with no connection ever made from their perspective, and the dialer has already moved on.

This is intentional behavior. The dialer is constantly protecting the agent pool from being overwhelmed by live contacts with no one to handle them. And rather than simply avoiding drops, the algorithm actively hunts for the highest possible contact rate it can sustain while keeping drops just under the configured threshold.

Session Capacity and the 375-Channel Ceiling

Maximum Research has 375 concurrent call channels provided by the telco. Think of these as 375 phone lines that can be active at the same time across all dialing activity.

However, in normal operations, the 375-channel ceiling is rarely the actual constraint. The real practical limit is the "3 calls per waiting agent" rule combined with the 30 calls per second cap. Both apply at all times.

Example:

- 30 agents on the floor x 3 calls each = 90 channels maximum
- 375 channels available from telco
- 285 channels sitting completely idle

This is by design. The 3:1 rule was set because at our peak staffing of 100 to 110 agents, the math works out to:

- 110 agents x 3 calls = 330 channels needed at peak
- 375 available from telco
- 45-channel buffer (about 14% headroom)

This ensures that even on our busiest possible night, we never hit the telco ceiling and create a bottleneck.

Why Fewer Agents Can Mean More Dials Per Agent

This is counterintuitive but important. When fewer agents are working, the drop rate risk is lower because fewer calls are connecting overall, which allows the algorithm to dial more aggressively per agent.

Example comparison:

Night A -- 38 agents, high connect rate (7%)

- $38 \times 3 = 114$ max calls at once
- High connect rate means agents are busy -- dialer does not need to push as hard
- Drop risk is lower, but so is the need to dial aggressively
- Result: roughly 326 calls per hour

Night B -- 29 agents, lower connect rate (6.4%)

- $29 \times 3 = 87$ max calls at once
- Low connect rate means agents are idle more often -- dialer pushes closer to 3:1
- Greater drop risk, but algorithm pushes hard to stay just under the target
- Result: roughly 350 calls per hour

Night B made more calls per hour despite fewer agents. The algorithm pushing closer to the 3:1 cap to compensate for the lower contact rate was the deciding factor.

How Multiple Jobs Share Channels

Each job calculates its own dialing needs independently. In most cases, multiple jobs running simultaneously have no meaningful impact on each other, because the combined agent count across all jobs multiplied by 3 stays well under the 375-channel ceiling.

Channel competition only becomes a real factor if total active agents across all running jobs would push the combined demand past 375. If that threshold were approached, jobs with more agents would carry more weight in claiming available channels, and the per-job drop rate and connect rate at that moment would further influence how many channels each job actually receives.

Under normal staffing conditions at Maximum Research, this ceiling is not a practical concern.

What This Means Operationally

- Early shift behavior is aggressive, not conservative. Because connect rate and drop rate values start at zero, the dialer has nothing to constrain it and immediately pushes toward the 3:1 cap. The result is that the majority of dropped calls typically occur early in a shift, with large swings between the actual drop rate and the target until enough call volume accumulates for the pattern to normalize.
 - Staffing level directly affects dialing behavior. More agents means more channels consumed, but a higher contact rate also means agents are busier and the dialer does not need to push as hard. There is a sweet spot -- typically around 25 to 35 agents for a single job -- where the algorithm can push closest to the drop threshold while maximizing contacts.
 - Two jobs running simultaneously is not a problem under normal staffing. Each job manages its own dialing independently and draws from the shared channel pool without conflict, as long as combined session demand stays under 375.
 - The 30 calls per second cap is a hard system limit that exists independently of everything else. On a typical night we do not come close to it, but it is there as a safeguard against runaway dialing.
-

A Note on the Dashboard

The Survox dialer dashboard updates its displayed statistics every 10 to 15 seconds. This is intentional -- frequent screen refreshes create overhead, and the dashboard is a reporting display, not the calculation engine. The dialer itself is recalculating and firing far faster than what the dashboard reflects. Do not assume that what you see on the dashboard represents the dialer's current state at that exact moment.

Summary

The Survox predictive dialer is a self-regulating system. It is event-driven -- each job recalculates independently every time its agents are waiting and more dials are needed -- rather than running on a fixed clock. Its goal is not simply to stay within the configured drop threshold, but to find the highest possible contact rate it can sustain while landing just below that threshold.

Because of this, the dialer is always running as fast as the current conditions allow for each job. If a job is already hitting the 3:1 capacity ceiling and the drop target is not being exceeded, raising the drop target will have no effect -- the dialer is already at maximum speed and the drop rate simply reflects the contact conditions that exist.

Multiple jobs running simultaneously operate independently without interfering with each other, unless combined session demand across all jobs approaches the 375-channel ceiling -- a situation that does not occur under normal staffing conditions at Maximum Research.

Revision #3

Created 2026-02-27 22:01:35 UTC by Admin

Updated 2026-03-04 16:44:43 UTC by Admin