replicant

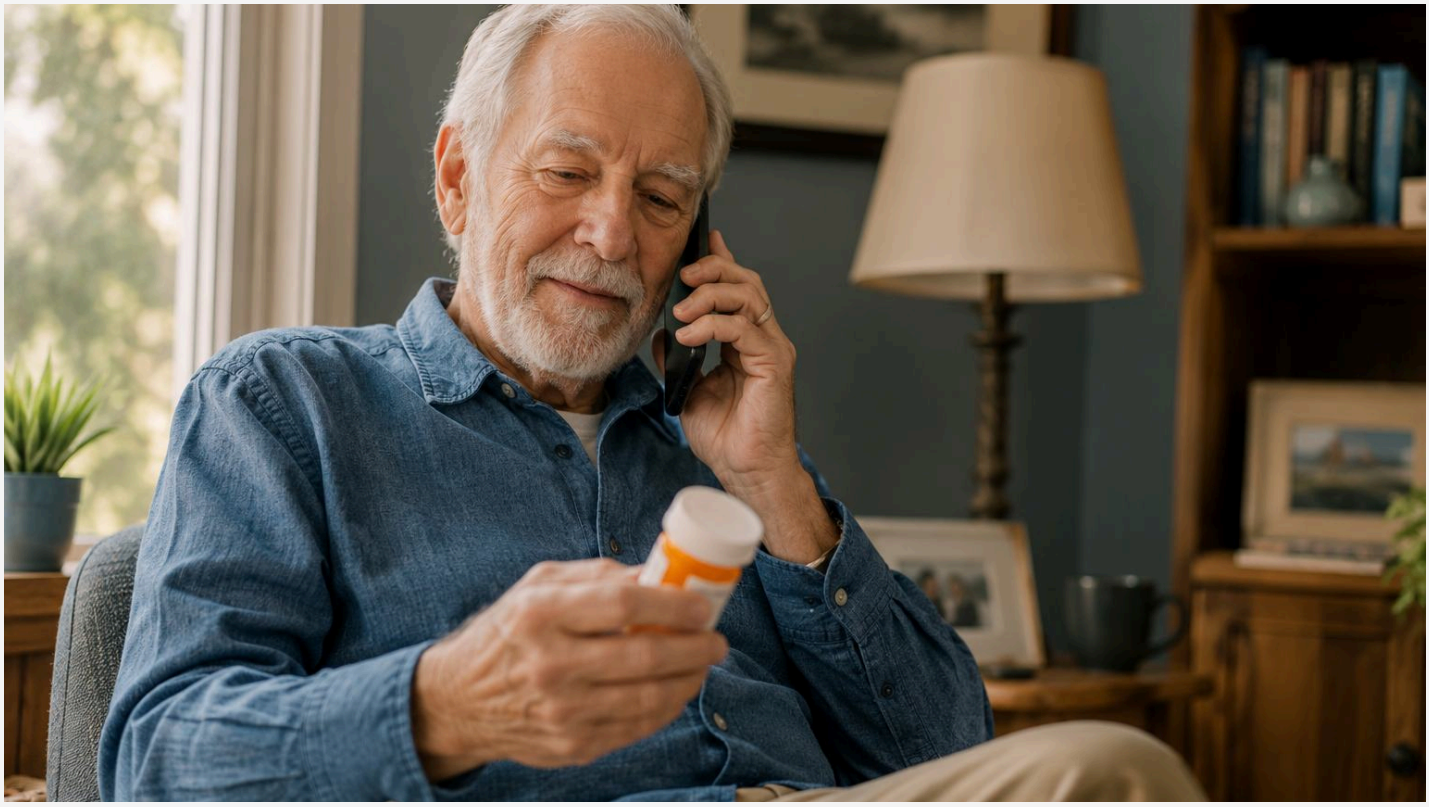

Healthcare

The Healthcare Phone Line

What roughly 460,000 analyzed voice calls, across six healthcare organizations spanning payers, providers, diagnostics, imaging, and pharmacy, reveal about demand, resolution, and where conversational AI creates value.

Anonymized industry insight | Six healthcare deployments | Trailing 12 months | July 2026

In Brief

1. **Healthcare voice is not one call type, it is two.** A high-volume transactional layer (scheduling, account, results) sits alongside long clinical and benefits conversations. Call length spans roughly **one minute to nineteen** across organizations, so a single benchmark misleads.

2. **The transactional layer is already automatable and resolves well.** Where scheduling leads, resolution reaches **~90%;** where calls are clinical or benefits-heavy, most end only **partially resolved.** Outcomes range from ~26% to ~90% resolved by organization.

3. **The prize is asymmetric.** Automate appointments, account servicing, and results end to end; augment clinical judgment, authorizations, and eligibility with a human in the loop rather than replacing them.

What these numbers measure

A diagnosis of the channel as it runs today, not Replicant's automation results

The figures in the sections below describe the **current phone channel** at these healthcare organizations (existing agents and IVR), measured by Replicant's Conversation Intelligence. Dispositions are **whole-channel**, not AI-only containment. This sizes the opportunity; **Replicant's own deployment results are in Section 5**, where containment (calls handled without a live agent) is reported separately.

Section 1

Healthcare voice splits into transactional and clinical

Across the six organizations, demand divides cleanly. One set of calls is transactional and self-contained, scheduling an appointment, checking an account or a result, finding a location. The other is clinical or benefits-oriented, a medical question, an authorization, an eligibility check, and runs far longer. The mix a given organization sees, more than any single benchmark, determines its call length, resolution, and where AI helps.

1-19min

Range of average call length across organizations

26-90%

Range of full resolution across organizations

~76%

Share of calls that are appointments at imaging-led providers

Section 2

Scheduling and account servicing are the transactional core

Grouping reasons into themes, the recurring, automatable drivers are appointments and scheduling, account and member servicing, benefits and eligibility, and results, orders, and delivery. Clinical questions and authorizations are a smaller but higher-complexity slice. The exact mix varies widely by organization, so shares below are directional across the cohort.

Exhibit 1

What healthcare consumers call about (directional, cohort-wide)

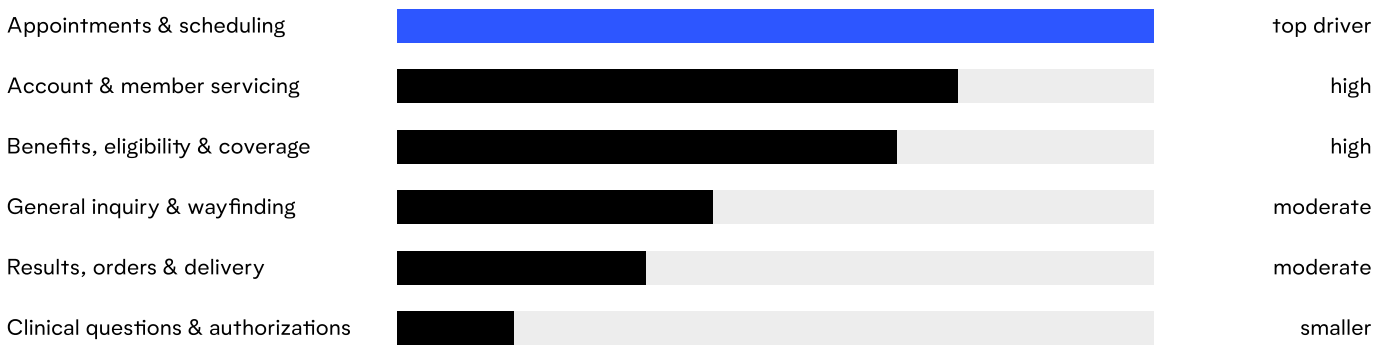

Bars show relative prominence across the cohort, not precise shares; reason taxonomies differ by organization (e.g. appointments reach ~76% of calls at an imaging provider but ~30% at a claims organization). Source: Replicant Conversation Intelligence, trailing 12 months.

Section 3

Outcomes vary more than in any other vertical

Equal-weighted across the six organizations, about 44 percent of calls fully resolve. But the dispersion is the headline: full resolution ranges from roughly a quarter to nearly all calls, abandonment from near zero to almost 60 percent, and transfers from a few percent to nearly half. The driver is call type. Scheduling-led programs resolve cleanly; clinical and benefits programs leave most calls partially resolved and lean on people.

Exhibit 2

How healthcare calls end (equal-weighted; ranges show the spread)

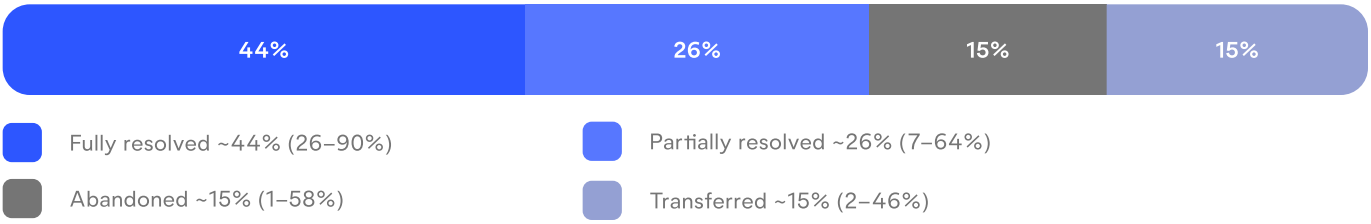

Share of calls by final disposition, equal-weighted across six organizations; ranges show organization minimum to maximum. Wide ranges reflect different call-type mixes, not measurement noise. Source: Replicant Conversation Intelligence, trailing 12 months.

Section 4

Where to focus first: a priority map

Positioning healthcare call types by their readiness for automation against volume and value points to a clear sequence. Appointments and account servicing sit top-right, high volume and ready to automate now. Benefits and eligibility bring value but need payer and system integration. Clinical questions and complex authorizations stay with clinicians, augmented rather than replaced.

Exhibit 3

Automation priority map: readiness versus volume and value

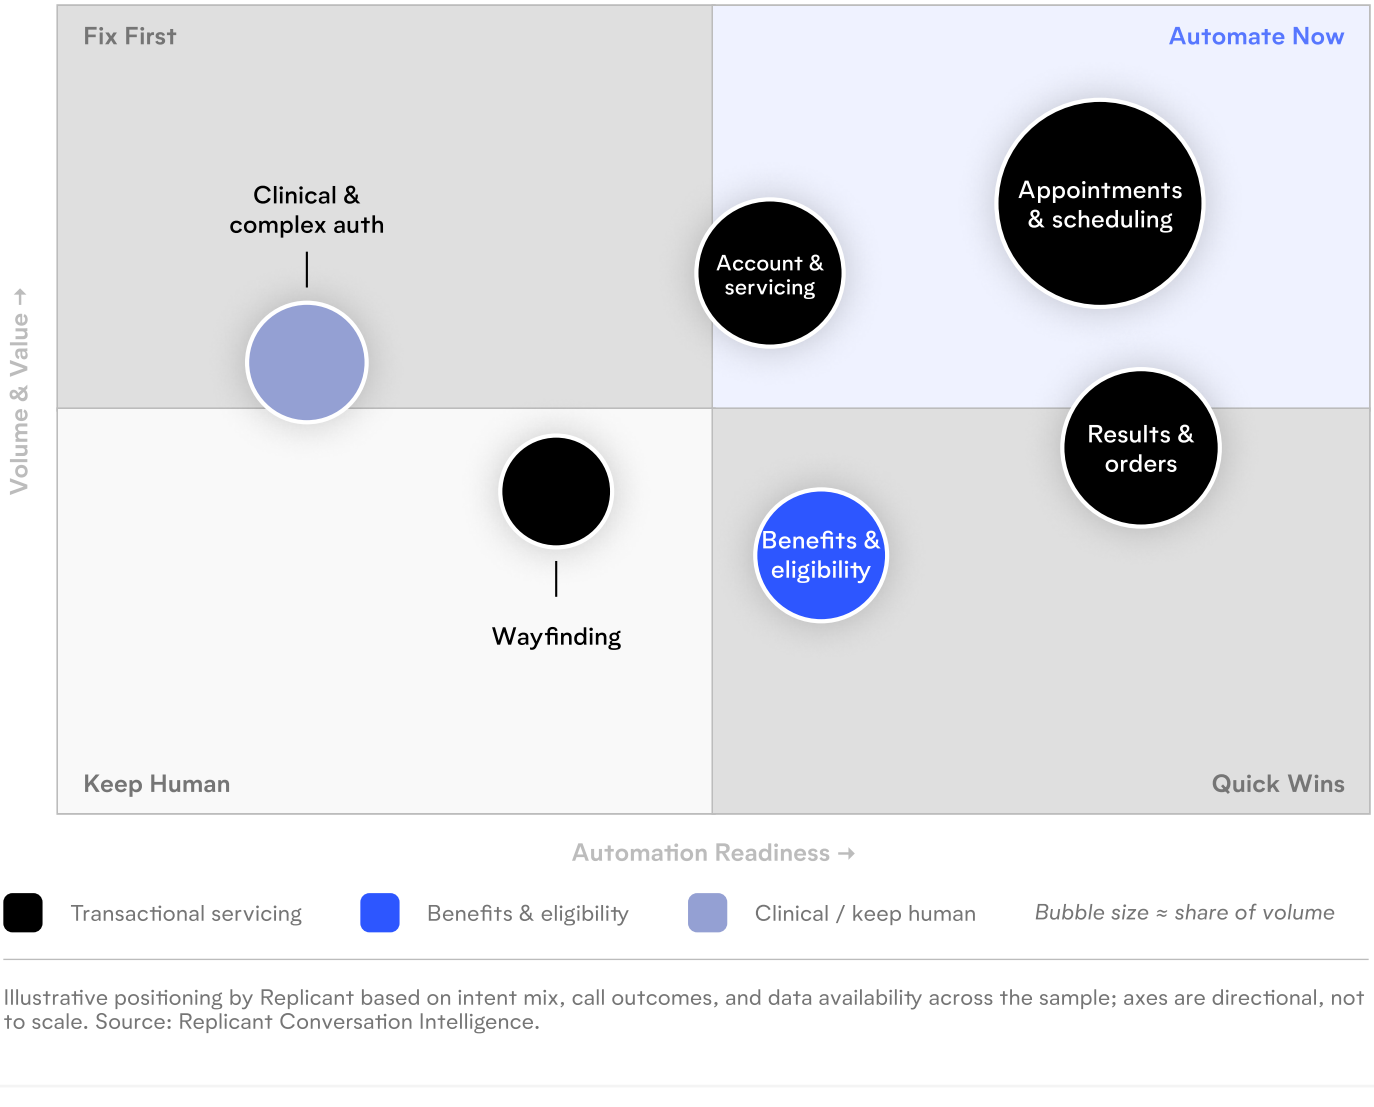

Illustrative positioning by Replicant based on intent mix, call outcomes, and data availability across the sample; axes are directional, not to scale. Source: Replicant Conversation Intelligence.

Section 5

From diagnosis to outcome: Replicant in healthcare

The sections above size the opportunity. Across Replicant's healthcare deployments, automation converts routine, transactional demand, scheduling, results, account, into measured containment and lower abandonment at scale, freeing staff for clinical and complex work. Containment here means calls handled without a live agent, distinct from the whole-channel resolution measured in the diagnostic.

700+/day

Calls automated at one provider

-20%

Call abandonment at that provider

500k+

Calls automated at a claims organization

Exhibit 4

Measured results from Replicant auto-club deployments

Deployment	Measured result
Medical imaging provider	**700+** calls/day automated; call abandonment **−20%**
Claims / care-management org	**500k+** calls automated; runs on Conversation Intelligence

Results from Replicant healthcare deployments, shown as anonymized descriptors and distinct from the diagnostic cohort. Containment / automation are deployment metrics, distinct from the whole-channel resolution measured in the diagnostic.

Section 6

Built for healthcare: HIPAA, security, and PHI protection

Healthcare voice touches PHI, consent, and clinical judgment. Replicant is built to operate inside that framework, with enterprise security certified by independent audit and conversation-level controls configured to each organization's requirements.

Security & Data privacy	Controls for healthcare conversations
Certified. SOC 2 Type 2 (independent annual audit), HIPAA, PCI, and GDPR compliant.	**PHI redaction.** Protected health information redacted from transcripts and recordings.
Encrypted. AES-256 at rest (GCP KMS); TLS 1.2 in transit.	**Consent & recording.** Configurable consent capture and recording disclosures.
Hosted in the US. Google Cloud, continental US; redundant, high-availability infrastructure.	**Human in the loop.** Guaranteed escalation for clinical and complex cases.
Access-controlled. Role-based access, IP restrictions, network segmentation, need-to-know.	**Consistent delivery.** Every caller hears the same approved script, verbatim.
Tested. Independent penetration testing, bug bounty, IPS/IDS monitoring, incident-response plan.	**Your framework.** Retention and access configured to your policies and regulations.

Platform security and data privacy per Replicant's Security & Data documentation. Conversation-level controls are configured to each organization's compliance program.

Section 7

Two agendas: automate the routine, augment the clinical

Automate (Transactional)	Augment (Clinical/Benefits)
Appointments & scheduling. Book, reschedule, cancel, and remind, 24/7.	**Benefits & eligibility.** Pre-check coverage and gather details; route the exception to a specialist.
Account & member servicing. Verify identity, update details, answer plan and account questions.	**Authorizations & claims.** Collect and validate; escalate the decision to a human.
Results, orders & delivery. Status of a test, order, or delivery on request.	**Clinical questions.** Triage and hand off with context; never replace clinical judgment.
Wayfinding & general info. Locations, hours, preparation instructions.	**Proactive outreach.** Reminders and follow-ups to cut inbound volume and no-shows.

The next step

This diagnosis is directional, six organizations, not your book. Replicant can run this same Conversation Intelligence analysis on **your** calls, quantify where your channel leaks, and model the containment and experience lift from automating it, before any commitment.

Methodology & scope

Market figures are drawn from Replicant Conversation Intelligence over the trailing 12 months and are equal-weighted across six production healthcare deployments spanning claims and care management, health and wellness benefits, diagnostics, home infusion, medical imaging, and a Medicare Advantage plan (~460,000 analyzed calls); each organization contributes equally and the organization minimum-to-maximum range is shown. Driver prominence is directional because reason taxonomies differ by organization. The priority map is an illustrative synthesis of intent mix, outcomes, and data availability; its axes are directional, not to scale. Automated quality scores use per-program rubrics and are not compared across organizations. No organization, client, patient, or individual caller is identified; no protected health information or underlying data left the analytics environment.

Skip the demo. Start building your AI roadmap.

Get these insights for your contact center with a no-cost Conversation Intelligence analysis.

Learn more