



Retail & E-Commerce

The Retail Phone Line

What roughly 2.4 million analyzed voice calls, across four retail and e-commerce brands spanning merchandise, apparel, and event ticketing, reveal about demand, resolution, and where conversational AI creates value.

Anonymized industry insight

Four brands

Trailing 12 months

July 2026

In Brief

- 1 Retail voice is the **post-purchase queue**. Order status ("where is my order"), returns, refunds, and product questions dominate demand, repetitive, structured, and highly automatable.
- 2 It carries the **lowest caller sentiment of any vertical** (~5.3 of 10). These are frustrated returns, refund, and missing-order calls, an experience risk at the exact moments that decide loyalty and repeat purchase.
- 3 Only about **two in five calls resolve** today. Automating the post-purchase queue, and absorbing seasonal peaks without adding staff, is the prize.

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 retail brands (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

Demand is post-purchase servicing

Grouping reasons into themes, retail demand concentrates in a handful of post-purchase intents: order status and tracking, returns and refunds, product and order support, and, for ticketing, order and ticket access. These are the most repeatable, most automatable requests in the contact center. Shares below are directional across the cohort, since taxonomies differ by brand.

Exhibit 1

What retail customers call about (directional, cohort-wide)



Bars show relative prominence across the cohort, not precise shares; reason taxonomies differ by brand. Source: Replicant Conversation Intelligence, trailing 12 months.

Section 2

The frustrated queue: lowest sentiment, half unresolved

Retail carries the lowest caller sentiment of any vertical we measure. That is unsurprising given the mix, no one enjoys chasing a refund or a missing order, but it is a loyalty risk precisely where retention is won or lost. Equal-weighted across the four brands, about 41 percent of calls fully resolve; the rest end partially resolved, abandoned, or transferred.

5.3/10

Caller sentiment, the lowest of any vertical (range 4.8–6.1)

Top 3


Drivers are all post-purchase: status, returns, refunds

~59%

Of calls do not fully resolve today

Exhibit 2

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



Share of calls by final disposition, equal-weighted across four brands; ranges show brand minimum to maximum. Source: Replicant Conversation Intelligence, trailing 12 months.

Section 3

Volume is spiky: the peak is the test

Retail demand is seasonal and event-driven, holidays, product drops, and, in ticketing, on-sales. The post-purchase queue swells exactly when staffing is hardest to flex, and abandonment and wait times rise with it. Because the demand is repetitive and structured, it is well suited to automation that scales instantly to absorb the peak, holding experience steady when volume triples rather than degrading at the worst possible moment.


Section 4

Where to focus first: a priority map

Positioning retail call types by their readiness for automation against volume and value points to a clear sequence. Order status and product support sit top-right, high volume and ready to automate now. Returns and refunds are high value but need policy and systems integration. Disputes and complex complaints stay with people, augmented. Ticket/account access stay with people, augmented. Delivery & shipping sit bottom-right, low volume and ready to automate now. Product & support sit bottom-right, low volume and ready to automate now. Quick Wins sit bottom-right, low volume and ready to automate now.

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 retail

The sections above size the leak. Across Replicant's retail and commerce deployments, automation converts the post-purchase queue into measured containment and higher resolution, absorbs peak surges without hiring, and lifts NPS on the calls that shape loyalty. Containment and resolution here are deployment outcomes, distinct from the whole-channel resolution measured in the diagnostic.

25→48%

AI containment before and after, at one retailer

+42

NPS improvement (21 to 63) at that retailer

8→22%

Resolution rate before and after

Exhibit 4

Measured results from Replicant retail & commerce deployments

Deployment	Measured result
Sports-merchandise retailer	Containment 25% → 48%; resolution 8% → 22%; NPS 21 → 63; 25k calls/day; hiring -40%
Large online marketplace	Resolution 87%; 350k calls/day automated; launched in 6 weeks

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

Section 6

Enterprise security and data privacy

Retail voice handles payments, refunds, and personal information at high volume. Replicant runs on an enterprise-secure platform certified by independent audit, with conversation-level controls configured to each brand's requirements.

Security & Data privacy

Certified. SOC 2 Type 2 (independent annual audit), HIPAA, PCI, and GDPR compliant.

Encrypted. AES-256 at rest (GCP KMS); TLS 1.2 in transit.

Hosted in the US. Google Cloud, continental US; redundant, high-availability infrastructure.

Access-controlled. Role-based access, IP restrictions, network segmentation, need-to-know.

Tested. Independent penetration testing, bug bounty, IPS/IDS monitoring, incident-response plan.

Controls for customer conversations

PCI-secure payments. Handle payments and refunds within a PCI-compliant platform.

Consent & recording. Configurable consent capture and recording disclosures.

PII redaction. Sensitive data redacted from transcripts and recordings.

Human in the loop. Guaranteed escalation for disputes and complex complaints.

Your framework. Retention and access configured to your policies.

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

Section 7

The playbook: clear the queue, protect the peak

Clear the queue

Order status & tracking. Answer "where is my order" instantly, 24/7.

Returns & refunds. Initiate, track, and status returns and refunds end to end.

Product & order support. Product questions, order changes, and cancellations.

Ticket / account access. Retrieve tickets and resolve access issues on the call.

Protect the peak & the experience

Scale instantly. Absorb holiday, drop, and on-sale surges without staffing.

Cut the wait. Eliminate hold and abandonment when it matters most.

Defuse frustration. Fast, accurate answers on the calls that shape loyalty.

Escalate with context. Hand disputes to a person with the full history.

The next step

This diagnosis is directional, four brands, not your book. Replicant can run this same Conversation Intelligence analysis on **your** calls, quantify where the post-purchase queue leaks, and model the containment and CSAT 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 four production retail and e-commerce brands spanning merchandise, apparel, fan commerce, and event ticketing (~2.4 million analyzed calls); each brand contributes equally and the brand minimum-to-maximum range is shown, so the largest brand does not drive a headline. Driver prominence is directional because reason taxonomies differ by brand. 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 brands. No brand, client, or individual caller is identified, and no 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](#)