

AI-FIRST HELP DESK MIGRATION

Protect your AI baseline. Migrate in the right order.

The two-step framework that trains your AI on clean data from day one
— and preserves your full historical record in the background.

96%

of orgs face data barriers
to AI adoption

The platform is new. The AI is already behind.

- Teams migrate and watch AI intent detection drop within the first week
- Suggested replies turn stale — knowledge base surfaces outdated articles
- Bulk-importing years of unresolved tickets distorts every AI model

Most teams move the data. Almost none move it in the right order.

< 40%

baseline triage accuracy
when AI enabled before
data cleanup

Twig, 2026

Your AI is only as good as the data feeding it.

Modern AI runs a continuous NLP training loop — scanning resolved tickets, extracting patterns from successful outcomes, applying them to classify new requests. Feed it bad data, it learns bad patterns.



Resolved tickets only

Open tickets have no outcome — they add noise, not signal.



CSAT \geq 4 stars

High ratings = your team's best practices, not the bad days.



Last 12–18 months

Older tickets teach outdated pricing and deprecated features.



Intent tags

Pre-labelled data lets AI map patterns significantly faster.



KB articles first

The KB is the primary retrieval source — migrate it before tickets.

AI-accurate data first. Full history second.

STEP 1

AI-Ready Migration

Move only what builds a clean AI baseline — resolved tickets filtered by recency and CSAT, all KB articles, and associated contacts.

AI working from week one

STEP 2

Full Historical Migration

All remaining tickets, contacts, companies, and attachments — migrated in chunks around business hours after the AI baseline is validated.

History with zero AI contamination

The sequence is fixed: Step 2 never begins until AI accuracy is validated at $\geq 85\%$.

Five filters that protect your AI baseline

1

Resolved & closed tickets only

Open tickets have no confirmed outcome — they add noise, not signal.

2

Last 12–18 months

Older tickets reference deprecated features, stale pricing, and policies.

3

CSAT score \geq 4 stars

High-rated interactions represent your team's best practices.

4

Active tags & intent categories

Structured tags give AI pre-labelled training data — skip if inconsistent.





5

All KB articles (every language)

Migrate KB before tickets — it's the primary AI retrieval source.

BEFORE YOU START

Disable on the target platform:

-  All outbound notifications
-  Live workflows & active triggers
-  Automated surveys
-  Any scheduled automations

Skipping this triggers automated surveys and duplicate replies mid-migration.

Don't flip channels until both thresholds are green.

Run a 50-ticket spot check using tickets excluded from Step 1. Two gates to pass.

GATE 1

$\geq 85\%$

AI intent detection accuracy
(industry standard — pause below 80%)

GATE 2

Zero

KB hallucinations — AI must not reference
content absent from your knowledge base

Zendesk, Intercom & Freshdesk each have different AI data rules.



Zendesk

Advanced AI / Intelligent Triage

- Migrate KB before a single ticket arrives
- Audit: remove duplicate & outdated articles
- Enable Intelligent Triage on one queue first
- Data first → triage accuracy starts at 40% without cleanup



Intercom

Fin AI Agent

- Fin reads KB + conversation history — both matter
- Prioritise logs with clear resolution + positive CSAT
- Published baseline: 65% resolution rate on curated data
- Unfiltered dump guarantees underperformance



Freshdesk

Freddy AI

- Auto-triage & article suggestions both need clean KB
- Growth plan minimum required for Freddy AI
- Same 5 filters apply — no exceptions
- Raw historical data corrupts macro recommendations

Then bring the complete record. With zero AI risk.

Once your team validates Step 1, migrate everything else — at your own pace.



Compliance & audit

GDPR, HIPAA & industry regulations mandate retaining support records. Leaving history on a lapsed source licence creates audit risk.



Agent context

Customers reference 2–3 year old tickets daily. Without history, agents rebuild context mid-conversation.



Reporting continuity

Missing data distorts trend analysis, volume reporting, and SLA tracking. Complete dataset = accurate analytics.

MONITOR ACCURACY

Watch intent detection after each chunk. Flag and pause if:

**Accuracy drops > 5%
from Step 1 baseline**

Audit the chunk before continuing — that drop means a batch of low-quality or very old tickets contaminated the model.



TRY IT ON YOUR OWN DATA

Run a free Demo migration before you commit

See exactly which tickets and KB articles qualify for Step 1.

Preview how your archive sequences in Step 2.

No commitment. No downtime. Zero data left behind.

[Start your free Demo migration →](#)

90+ platforms

GDPR

HIPAA

SOC 2