

# TPPC Review: High-risk support escalation approval

This workflow is a good candidate for a governed pilot because the breakage is concentrated in one repeatable step: agents produce weak escalation writeups, and managers spend time reconstructing context before they can approve or deny an exception. The human boundary is clear and should remain intact: managers must keep authority over refund exceptions, policy deviations, and the final customer-facing decision. The narrow AI opportunity is to draft a structured escalation summary for the manager using the existing ticket context and policy references. That targets delay and rewrite volume without automating the approval itself.

### THE DECISION

**Proceed with a narrow pilot focused only on drafting the escalation brief.**

**Proceed to governed pilot / 88% confidence**

## Northwind Support Operations / High-risk support escalation approval

Prepared automatically 2026-08-10 | Secure decision support - not a guarantee of outcomes

ILLUSTRATIVE SAMPLE | Fictional company and workflow data | Not a customer result

# TPPC Review: High-risk support escalation approval

## EXECUTIVE DECISION

**Proceed with a narrow pilot focused only on drafting the escalation brief.**

The workflow has a clear owner, a defined trigger, measurable pain, a preserved human approval boundary, and usable pilot evidence. The main waste is not the decision itself; it is the repeated reconstruction and rewriting before the decision. A draft-first assist is narrow enough to test safely and directly aligned to the stated goals.

|                                                     |                                 |                                                   |
|-----------------------------------------------------|---------------------------------|---------------------------------------------------|
| <b>DECISION</b><br><b>Proceed to governed pilot</b> | <b>CONFIDENCE</b><br><b>88%</b> | <b>OWNER</b><br><b>Support Operations Manager</b> |
|-----------------------------------------------------|---------------------------------|---------------------------------------------------|

## CURRENT-STATE BRIEF

**Agents summarize the case manually, duplicate the summary into Slack, wait for a manager, then managers reconstruct context, rewrite the escalation, and record the final decision back in Zendesk.**

## OPERATING FACTS

- Trigger: A frontline agent flags a refund, policy, or customer-risk exception in Zendesk.
- Objective: Reduce median approval time below four hours and reduce manager rewrites by 40%, while maintaining zero unapproved policy exceptions.
- Systems: Zendesk, Slack, Salesforce knowledge base

# Where AI helps - and where it stops

## CURRENT TO TARGET FLOW

|    |                                                                                        |                                       |
|----|----------------------------------------------------------------------------------------|---------------------------------------|
| 01 | Frontline agent identifies a refund, policy, or customer-risk exception in Zendesk.    | <b>HUMAN</b><br>Frontline agent       |
| 02 | Agent prepares the escalation summary for manager review.                              | <b>AI ASSISTED</b><br>Frontline agent |
| 03 | Agent posts the escalation into Slack for manager attention.                           | <b>HUMAN</b><br>Frontline agent       |
| 04 | Manager reviews the case, checks policy context, and approves or denies the exception. | <b>HUMAN</b><br>Support manager       |
| 05 | Manager or agent records the final rationale and outcome back in Zendesk.              | <b>HUMAN</b><br>Support manager       |
| 06 | Agent sends the final customer-facing response after manager approval.                 | <b>HUMAN</b><br>Frontline agent       |

## RECOMMENDED AI-OWNED ASSISTANCE

### DRAFT: Agent prepares the escalation summary for manager review.

This is where delay and rework start. Managers are rewriting most escalations because the initial brief is incomplete or inconsistent. A structured draft can improve completeness and readability before the case reaches the manager, without shifting approval authority. Expected gain: Faster manager review, fewer manager rewrites, and more consistent rationale capture across Zendesk and Slack.

## PROHIBITED AI ACTIONS

- Do not approve or deny any refund exception or policy deviation.
- Do not generate or send the final customer-facing decision without manager approval.
- Do not omit the source ticket link, evidence, or cited policy reference when available.

# Human authority, fallback, and material risks

## HUMAN AUTHORITY

**The support manager must make the final approval decision on refund exceptions, policy deviations, and any final customer-facing outcome.**

## APPROVAL TRIGGER

**Every flagged high-risk escalation requires explicit manager approval before any exception is granted or any final response is sent.**

## SAFE FALLBACK

**If the AI draft is missing facts, cites the wrong policy, or is not trusted by the agent or manager, the team reverts to the current manual summary process for that case and records the issue for prompt or template correction.**

Use this path whenever confidence, policy, or data quality drops below the agreed threshold.

## AUDIT EVIDENCE

**Zendesk ticket ID and timestamps; Original agent draft or AI-assisted draft; Slack escalation message link; Manager edits to the draft; Final approval or denial outcome; Recorded rationale in Zendesk.**

# What could fail and how the pilot contains it

## MATERIAL RISKS

|        |                                                                                                                     |                                                                                                                                                                                                              |
|--------|---------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HIGH   | <b>The draft may misstate facts or omit a critical case detail, causing the manager to review a flawed summary.</b> | Use a fixed escalation template with required fields, require the agent to verify the draft before posting, and require the manager to approve based on the underlying ticket rather than the summary alone. |
| HIGH   | <b>The draft may cite the wrong policy or present policy guidance too confidently.</b>                              | Limit the AI output to a draft summary plus clearly labeled policy reference fields, and require the manager to confirm the policy basis before approval.                                                    |
| MEDIUM | <b>Slack duplication may still create version drift between the ticket and the escalation thread.</b>               | Treat Zendesk as the system of record during the pilot and require the Slack message to include the Zendesk ticket link and the same structured summary.                                                     |
| MEDIUM | <b>Medium-sensitivity customer data may be exposed more broadly during testing.</b>                                 | Use redacted historical cases for offline evaluation, restrict pilot access to the existing support team, and avoid expanding data sharing beyond Zendesk and Slack.                                         |
| MEDIUM | <b>Managers may continue rewriting out of habit, hiding whether the draft actually improves quality.</b>            | Track manager edits by category, review examples weekly, and tune the template against the most common rewrite reasons.                                                                                      |

## ASSUMPTIONS TO VERIFY

- The team can use 30-50 redacted escalations with timestamps, manager edits, approval outcomes, and policy-incident counts to establish and compare pilot performance.
- A pilot can begin without changing approval authority or replacing Zendesk and Slack.
- The Salesforce knowledge base contains relevant policy material, but direct retrieval quality is not yet proven and should be validated during offline testing.
- The team can measure manager rewrite volume from draft-to-final comparisons in the sample and pilot cohort.

# Prove value before expanding scope

- 1** **WEEK 1** Pull the 30-50 redacted escalation sample from Zendesk and Slack. Define one required escalation template for the pilot: issue, customer impact, requested exception, evidence, policy reference, risk note, and recommended manager review point. Confirm Zendesk remains the system of record and document the current approval timestamp and rewrite baseline from the sample.
- 2** **WEEK 2** Run an offline evaluation on the redacted sample. Use the AI only to draft the structured escalation brief from the existing case context. Score each draft for completeness, factual accuracy against the ticket, and amount of manager rewriting needed. Keep a defect log for missing facts, wrong policy references, and unclear rationale.
- 3** **WEEK 3** Start a live shadow pilot on a limited set of new high-risk escalations in the existing workflow. Agents review the AI draft, post the approved draft into Slack, and managers continue normal approval. Record approval time, manager edits, approval outcome, and any policy incidents for the pilot cohort.
- 4** **WEEK 4** Review pilot results with the Support Operations Manager and manager approvers. Keep the pilot only if the cohort shows median approval time below four hours, manager rewrites reduced by 40% versus baseline, and zero unapproved policy exceptions. If not, revert fully to the manual summary process and use the defect log to decide whether to redesign or stop.

## SUCCESS GATE

**For the pilot cohort, median approval time is below four hours, manager rewrites are reduced by 40% against the measured baseline, and there are zero unapproved policy exceptions.**

Proceed beyond pilot only if the AI-assisted draft cohort meets the stated success gate while preserving the manager approval boundary on every case.

# Measurement, rollback, and the next move

## ROLLBACK TRIGGER

**Any case of approval bypass, any unapproved policy exception, repeated factual errors in the draft that managers judge unsafe, or no measurable reduction in manager rewrite burden by the end of the pilot.**

If the AI draft is missing facts, cites the wrong policy, or is not trusted by the agent or manager, the team reverts to the current manual summary process for that case and records the issue for prompt or template correction.

## MEASUREMENT CONTRACT

**Median manager approval time for high-risk escalations.**

Use the provided 30-50 redacted escalations to measure current approval time, manager rewrite frequency, approval outcomes, and policy-incident counts before the live pilot.

## VALUE INTERPRETATION

**The value comes from reducing manager reconstruction and rewrite work, not from automating high-risk approvals.**

## NEXT BEST ACTION

**Approve a two-week offline evaluation using the redacted escalation sample, with one structured AI-drafted escalation brief as the only pilot scope.**