AI READINESS REPORT

acme-it.freshservice.com

# 53/100

**Partially Ready**

Your Freshservice instance has a solid SLA configuration but significant gaps in Knowledge Base and Service Catalogue are limiting Freddy AI to roughly a third of its potential. Fix the KB first.

| AI VALUE CAPTURE | EST. REMEDIATION | ISSUES IDENTIFIED | DIMENSIONS SCORED |
|---|---|---|---|
| **30%** | **72 hours** | **21** | **7 of 7** |

| 53/100 | 30% | 72h |
|---|---|---|
| OVERALL READINESS SCORE | AI VALUE CAPTURE TODAY | ESTIMATED REMEDIATION |

REPORT SUMMARY

## Partially Ready

AI features will partially work but miss significant value. Remediation is recommended before full rollout.

**FREE REPORT LAYER**

This PDF includes the score, dimension breakdown, top quick wins, and visible issue list. Blueprint-only fixes, paste-ready templates, and the full implementation plan are omitted.

**PRIORITY ACTION — DO THIS FIRST**

Publish at least 40 KB articles across your top 8 ticket categories before enabling Freddy AI chatbot — without content, the chatbot cannot answer questions and will deflect zero tickets.

YOUR INSTANCE AT A GLANCE

| 53/100 | 30% | 72h | Partially Ready |
|---|---|---|---|
| Overall score | AI value capture | Est. remediation | Readiness band |

## WHAT THIS SCORE MEANS

| | | |
|---|---|---|
| 90–100 | **Optimised** | Best-in-class. AI features will perform at or near their maximum potential. |
| 70–89 | **AI-Ready** | Solid configuration. AI features will work well with a focus on optimisation. |
| 40–69 | **Partially Ready • Your score** | AI will partially work but miss a significant portion of its potential. Remediation recommended. |
| 0–39 | **Not Ready** | Significant gaps. AI deployment without remediation will likely underperform or fail. |

## SEVEN AI-READINESS DIMENSIONS

Each dimension is weighted by its impact on AI feature performance. Scores below 70 indicate areas that should be remediated before full AI rollout.

| DIMENSION | WEIGHT | SCORE |
|---|---|---|
| **Knowledge Base** | 22% | 32 |
| **Ticket Data Quality** | 18% | 56 |
| **Categorisation** | 15% | 61 |
| **Canned Responses** | 10% | 44 |
| **SLA Configuration** | 13% | 84 |
| **Service Catalogue** | 12% | 38 |
| **Ticket Routing** | 10% | 71 |

## QUICK WINS — LOW EFFORT, HIGH IMPACT

1. Create KB articles for your top 5 ticket categories (under 4 hours each)
2. Remove 8 duplicate canned responses that are training AI on contradictory content
3. Add subcategories to all 6 categories — currently none have subcategories configured
4. Add approval workflows to your top 3 service catalogue items

# Knowledge Base

**32**/100

22% WEIGHT IN OVERALL SCORE

> Critical gap. With only 12 published articles and 67% stale content, Freddy AI chatbot cannot deflect tickets effectively. This is the highest-priority fix for AI readiness.

## 3 ISSUES IDENTIFIED

ISSUE 1

### Only 12 published articles across all categories

AI IMPACT

Freddy AI chatbot requires a minimum of 25 articles to generate useful responses. At 12 articles, deflection rate will be near zero regardless of how well the chatbot is configured.

RECOMMENDED FIX

Target 5+ articles per top ticket category. Start with Password Reset, Software Access, VPN Issues, Hardware Requests, and Onboarding — these typically account for 60–70% of ticket volume.

HIGH EFFORT

ISSUE 2

### 67% of articles have not been updated in over 12 months

AI IMPACT

Freddy AI surfaces stale content with the same confidence as current content. Outdated articles create false resolutions and reduce agent trust in AI suggestions.

RECOMMENDED FIX

Schedule a KB review sprint. Prioritise articles linked to current products and processes. Archive anything referencing retired systems or discontinued processes.

MEDIUM EFFORT

ISSUE 3

### Average article length is 89 words — below the 150-word minimum for effective AI retrieval

AI IMPACT

Short articles lack context needed for AI to accurately match tickets to solutions. Freddy AI performs significantly better with articles that include symptoms, root causes, and step-by-step fixes.

RECOMMENDED FIX

Expand the shortest 50% of articles. Add symptom descriptions, prerequisites, and resolution steps. Aim for 200–400 words per article with a clear resolution procedure.

MEDIUM EFFORT

# Ticket Data Quality

56/100

18% WEIGHT IN OVERALL SCORE

> Moderate. Subject lines are reasonable, but 34% of tickets arrive with generic subjects and 18% have no description at all — limiting the signal Freddy AI and auto-triage can learn from.

## 3 ISSUES IDENTIFIED

ISSUE 1

### 34% of tickets have generic subjects such as "Help", "Issue", or "Request"

AI IMPACT

Auto-triage classifies primarily from the subject line. Generic subjects force tickets into a catch-all group and lower Freddy's routing and priority-prediction accuracy.

RECOMMENDED FIX

Add a subject-line validation rule to the portal form requiring at least 10 characters, and show example subjects per request type to nudge requesters toward specifics.

LOW EFFORT

ISSUE 2

### Average description length is 22 words — below the 40-word threshold for reliable AI context

AI IMPACT

Freddy AI needs enough detail to match a ticket to a solution. Thin descriptions reduce answer confidence and increase the rate at which agents ignore AI suggestions.

RECOMMENDED FIX

Add guided fields (what you expected, what happened, steps already tried) to your top 3 ticket forms so requesters supply structured detail by default.

MEDIUM EFFORT

ISSUE 3

### 18% of tickets arrive with no description at all

AI IMPACT

Empty descriptions give the AI nothing to work with — these tickets cannot be triaged or deflected and always require manual agent follow-up.

RECOMMENDED FIX

Make the description field mandatory on the requester portal and email-to-ticket intake, with a minimum length so one-word submissions are rejected.

LOW EFFORT

# Categorisation

**61**/100

15% WEIGHT IN OVERALL SCORE

> Moderate. Categories exist but lack depth. 24% of tickets arrive without subcategories, limiting auto-triage accuracy. Adding subcategories across all 6 categories would materially improve AI performance.

## 3 ISSUES IDENTIFIED

ISSUE 1

### 24% of tickets in the last 90 days have no subcategory assigned

AI IMPACT

Freddy AI auto-triage trains on category/subcategory pairs. Tickets without subcategories dilute the training signal and reduce auto-assignment accuracy by an estimated 18–30%.

RECOMMENDED FIX

Create 3–5 subcategories per category based on your most common ticket titles. Enable mandatory subcategory selection for new tickets via field validation rules.

LOW EFFORT

ISSUE 2

### None of your 6 categories have subcategories configured

AI IMPACT

Without subcategories, Freddy auto-triage can only classify at the top-level category. This creates overly broad routing that ignores specialist group assignments.

RECOMMENDED FIX

Start with your two highest-volume categories. Define 3–5 subcategories per category, map each to a specific support group, and run for 30 days before evaluating auto-triage accuracy.

LOW EFFORT

ISSUE 3

### 2 categories have had zero tickets in the last 180 days

AI IMPACT

Unused categories pollute the dropdown and may confuse requesters, leading to miscategorised tickets that degrade AI training data over time.

RECOMMENDED FIX

Deactivate the 2 unused categories after confirming with the relevant teams. This reduces categorisation noise immediately.

LOW EFFORT

# Canned Responses

**44**/100

10% WEIGHT IN OVERALL SCORE

> Below average. 9 canned responses is insufficient for Freddy Copilot to provide reliable suggestions. 3 duplicate pairs are reducing quality. Expanding the library would immediately improve suggestion relevance.

## 3 ISSUES IDENTIFIED

ISSUE 1

### Only 9 canned responses — below the 20-response threshold for effective AI suggestions

AI IMPACT

Freddy Copilot uses your canned response library to generate reply suggestions. At 9 responses, coverage is too narrow, leading agents to ignore suggestions entirely.

RECOMMENDED FIX

Create responses for your 15 most common ticket scenarios: acknowledgement, information request, resolution confirmation, escalation notice, and closure templates across your top 3 categories.

LOW EFFORT

ISSUE 2

### 3 duplicate canned response pairs detected with >85% content similarity

AI IMPACT

Duplicate responses train Freddy on contradictory content, reducing confidence scores and degrading suggestion quality.

RECOMMENDED FIX

Review and merge the 3 duplicate pairs: "Password Reset Acknowledgement" / "Account Access Confirmation", "Request Received" / "Ticket Logged", and "Auto-Close Notice" / "Closing — No Response".

LOW EFFORT

ISSUE 3

### No canned responses use dynamic variables ({{requester_name}}, {{ticket_id}})

AI IMPACT

Freddy Copilot personalisation depends on dynamic field references. Static-only responses make AI suggestions feel generic and reduce agent adoption by approximately 25%.

RECOMMENDED FIX

Update your 5 most-used responses to include {{requester_name}} and {{ticket_id}}. Add {{assigned_group}} to routing notifications.

LOW EFFORT

# SLA Configuration

**84**/100

13% WEIGHT IN OVERALL SCORE

> Strong. 3 SLA policies cover all 4 priority levels with business hours correctly configured. Breach rates are within acceptable limits. Minor optimisation possible on response SLAs for low-priority tickets.

## 3 ISSUES IDENTIFIED

ISSUE 1

### Low-priority SLA response time of 24 hours is driving a 26% breach rate for that tier

AI IMPACT

High breach rates cause Freddy AI to misclassify priority. The AI learns from historical SLA adherence — chronic breaches skew its urgency assessment.

RECOMMENDED FIX

Extend low-priority response SLA to 36 hours or add a triage buffer. Alternatively, route all low-priority tickets to a shared queue with a 4-hour triage window before the response clock starts.

MEDIUM EFFORT

ISSUE 2

### SLA policies do not have automatic escalation notifications configured

AI IMPACT

Without escalation rules, Freddy cannot learn from escalation patterns. AI-based priority suggestions miss the urgency signals that manual escalation would otherwise capture.

RECOMMENDED FIX

Configure escalation rules at 75% and 90% of SLA window for all priority levels. Route escalation notifications to the relevant group manager.

LOW EFFORT

ISSUE 3

### Public holiday exclusions are not configured in business hours settings

AI IMPACT

Without public holiday exclusions, SLA timers run during non-working periods, artificially inflating breach rates and teaching Freddy AI that urgent tickets are lower priority than they actually are.

RECOMMENDED FIX

Update Business Hours settings to add 2025 public holidays for your locale. Add an out-of-hours auto-reply and hold SLAs during extended outages.

LOW EFFORT

# Service Catalogue

**38**/100

12% WEIGHT IN OVERALL SCORE

> Significant gap. 4 catalogue items with no categories and no approval workflows limits AI self-service capability. Freddy's self-service features depend on a rich catalogue with clear workflows.

## 3 ISSUES IDENTIFIED

ISSUE 1

### Only 4 service catalogue items — below the 10-item threshold for meaningful AI self-service

AI IMPACT

Freddy self-service can only offer what is in the catalogue. At 4 items, most requesters will not find what they need and will raise a ticket instead, eliminating self-service benefit.

RECOMMENDED FIX

Identify your top 10 request types from the last 90 days of ticket data. Convert the top 8 into catalogue items. Prioritise: software access, hardware, new user setup, VPN, password reset.

HIGH EFFORT

ISSUE 2

### Zero service catalogue items have approval workflows configured

AI IMPACT

Without approval workflows, Freddy cannot automate request fulfilment — all requests still require manual agent handling. This negates the primary ROI of AI self-service.

RECOMMENDED FIX

Add a simple manager approval workflow to your top 3 catalogue items. Start with Software Access, Hardware Request, and New User Setup.

MEDIUM EFFORT

ISSUE 3

### Service catalogue has no categories — all 4 items appear in a single uncategorised list

AI IMPACT

Without categories, Freddy cannot intelligently suggest catalogue items based on ticket context or history.

RECOMMENDED FIX

Create 3 catalogue categories: "Access & Accounts", "Hardware & Equipment", and "New Starter". Assign existing items to their relevant category before adding new items.

LOW EFFORT

# Ticket Routing

71/100

10% WEIGHT IN OVERALL SCORE

> Good foundation. 67% auto-assignment rate is above average. 4 groups are well-defined. Reassignment rate of 15% is within acceptable bounds but could be reduced with refined routing rules.

## 3 ISSUES IDENTIFIED

ISSUE 1

### 33% of tickets are manually assigned — above the target threshold of 25%

AI IMPACT

High manual assignment rates indicate routing rules are incomplete. Freddy AI routing uses historical assignment patterns to learn — gaps in automation create gaps in AI routing accuracy.

RECOMMENDED FIX

Review the 33% of manually assigned tickets from the last 30 days. Identify the top 3 ticket categories not covered by existing automation rules and create routing rules for each.

MEDIUM EFFORT

ISSUE 2

### No routing rules handle after-hours ticket assignment

AI IMPACT

Tickets raised outside business hours with no routing rule default to unassigned, creating a morning backlog that skews SLA metrics and corrupts Freddy's priority learning data.

RECOMMENDED FIX

Create an out-of-hours routing rule that assigns all P1 tickets to an on-call group and queues P2–P4 to the relevant group for next-morning triage.

LOW EFFORT

ISSUE 3

### Routing rules do not account for requester department or location

AI IMPACT

Department-based routing significantly improves auto-triage accuracy. Freddy AI learns faster with location-specific patterns.

RECOMMENDED FIX

Add department-based routing for your largest 3 departments. If departments are not captured on tickets, add a mandatory Department field to your main ticket form first.

MEDIUM EFFORT

## VISIBLE ISSUE SEQUENCE

This free report lists the visible issues available before Blueprint unlock. Specific fixes, templates, and the full implementation plan are omitted from the free PDF.

1. **34% of tickets have generic subjects such as "Help", "Issue", or "Request"** — LOW EFFORT
   Ticket Data Quality
   Add a subject-line validation rule to the portal form requiring at least 10 characters, and show example subjects per request type to nudge requesters toward specifics.

2. **18% of tickets arrive with no description at all** — LOW EFFORT
   Ticket Data Quality
   Make the description field mandatory on the requester portal and email-to-ticket intake, with a minimum length so one-word submissions are rejected.

3. **24% of tickets in the last 90 days have no subcategory assigned** — LOW EFFORT
   Categorisation
   Create 3–5 subcategories per category based on your most common ticket titles. Enable mandatory subcategory selection for new tickets via field validation rules.

4. **None of your 6 categories have subcategories configured** — LOW EFFORT
   Categorisation
   Start with your two highest-volume categories. Define 3–5 subcategories per category, map each to a specific support group, and run for 30 days before evaluating auto-triage accuracy.

5. **2 categories have had zero tickets in the last 180 days** — LOW EFFORT
   Categorisation
   Deactivate the 2 unused categories after confirming with the relevant teams. This reduces categorisation noise immediately.

6. **Only 9 canned responses — below the 20-response threshold for effective AI suggestions** — LOW EFFORT
   Canned Responses
   Create responses for your 15 most common ticket scenarios: acknowledgement, information request, resolution confirmation, escalation notice, and closure templates across your top 3 categories.

7. **3 duplicate canned response pairs detected with >85% content similarity** — LOW EFFORT
   Canned Responses
   Review and merge the 3 duplicate pairs: "Password Reset Acknowledgement" / "Account Access Confirmation", "Request Received" / "Ticket Logged", and "Auto-Close Notice" / "Closing — No Response".

8. **No canned responses use dynamic variables ({{requester_name}}, {{ticket_id}})** — LOW EFFORT
   Canned Responses
   Update your 5 most-used responses to include {{requester_name}} and {{ticket_id}}. Add {{assigned_group}} to routing notifications.

9. **SLA policies do not have automatic escalation notifications configured** — LOW EFFORT
   SLA Configuration

Configure escalation rules at 75% and 90% of SLA window for all priority levels. Route escalation notifications to the relevant group manager.

10

**Public holiday exclusions are not configured in business hours settings**

SLA Configuration

Update Business Hours settings to add 2025 public holidays for your locale. Add an out-of-hours auto-reply and hold SLAs during extended outages.

LOW EFFORT

11

**Service catalogue has no categories — all 4 items appear in a single uncategorised list**

Service Catalogue

Create 3 catalogue categories: "Access & Accounts", "Hardware & Equipment", and "New Starter". Assign existing items to their relevant category before adding new items.

LOW EFFORT

12

**No routing rules handle after-hours ticket assignment**

Ticket Routing

Create an out-of-hours routing rule that assigns all P1 tickets to an on-call group and queues P2–P4 to the relevant group for next-morning triage.

LOW EFFORT

METHODOLOGY

# How your score was calculated

The AI Readiness Audit connects directly to your Freshservice instance via the platform API using read-only credentials. Raw API keys are not stored. The audit reads your live configuration and 90 days of ticket patterns, stores computed metrics and report outputs for your account, and scores seven dimensions that determine how effectively AI features will perform.

## SCORING DIMENSIONS

| | | |
|---|---|---|
| **Knowledge Base** | 22% weight | Critical gap. With only 12 published articles and 67% stale content, Freddy AI chatbot cannot deflect tickets effectively. This is the highest-priority fix for AI readiness. |
| **Ticket Data Quality** | 18% weight | Moderate. Subject lines are reasonable, but 34% of tickets arrive with generic subjects and 18% have no description at all — limiting the signal Freddy AI and auto-triage can learn from. |
| **Categorisation** | 15% weight | Moderate. Categories exist but lack depth. 24% of tickets arrive without subcategories, limiting auto-triage accuracy. Adding subcategories across all 6 categories would materially improve AI performance. |
| **Canned Responses** | 10% weight | Below average. 9 canned responses is insufficient for Freddy Copilot to provide reliable suggestions. 3 duplicate pairs are reducing quality. Expanding the library would immediately improve suggestion relevance. |
| **SLA Configuration** | 13% weight | Strong. 3 SLA policies cover all 4 priority levels with business hours correctly configured. Breach rates are within acceptable limits. Minor optimisation possible on response SLAs for low-priority tickets. |
| **Service Catalogue** | 12% weight | Significant gap. 4 catalogue items with no categories and no approval workflows limits AI self-service capability. Freddy's self-service features depend on a rich catalogue with clear workflows. |
| **Ticket Routing** | 10% weight | Good foundation. 67% auto-assignment rate is above average. 4 groups are well-defined. Reassignment rate of 15% is within acceptable bounds but could be reduced with refined routing rules. |

## EFFORT ESTIMATES

Each issue carries an effort estimate reflecting the typical implementation time for a single person with Freshservice admin access. These are directional estimates, not project plans.

**Low effort**
**~4 hours**

**Medium effort**
**~16 hours**

**High effort**
**~40 hours**

## CONTACT

Questions about this report or need hands-on implementation support?

**hello@servicedeskbuilder.com**

servicedeskbuilder.com

This report was generated by Service Desk Builder's AI Readiness Audit on 1 May 2025. The audit reads your platform configuration and ticket patterns using read-only API access. Scores and recommendations are based on configuration best practices for AI feature deployment and are not a guarantee of AI performance outcomes. Re-run the audit after implementing fixes to track improvement.