

EXECUTIVE DECK COMPANION

# Summit Comfort HVAC

AI/Automation Audit — QA Preview

---

## Recommended First Move

Pilot CSR "Morning Rush" Minimum Data Set for 30 days.

## Buyer / Contact

Maria Chen — General Manager


## Audit Goal

Find AI/automation improvements for fragile dispatch + inconsistent CSR intake without replacing existing systems.

## Report Status

QA Preview

# How To Use This Deck

 This deck summarizes the recommended sequence and decision logic. Use the web report or Markdown source for full evidence, risks, and implementation detail.

## Executive Orientation


Use this deck for executive orientation — not as the full web report.

## Full Evidence

Use the web report or Markdown source for full evidence and implementation detail.

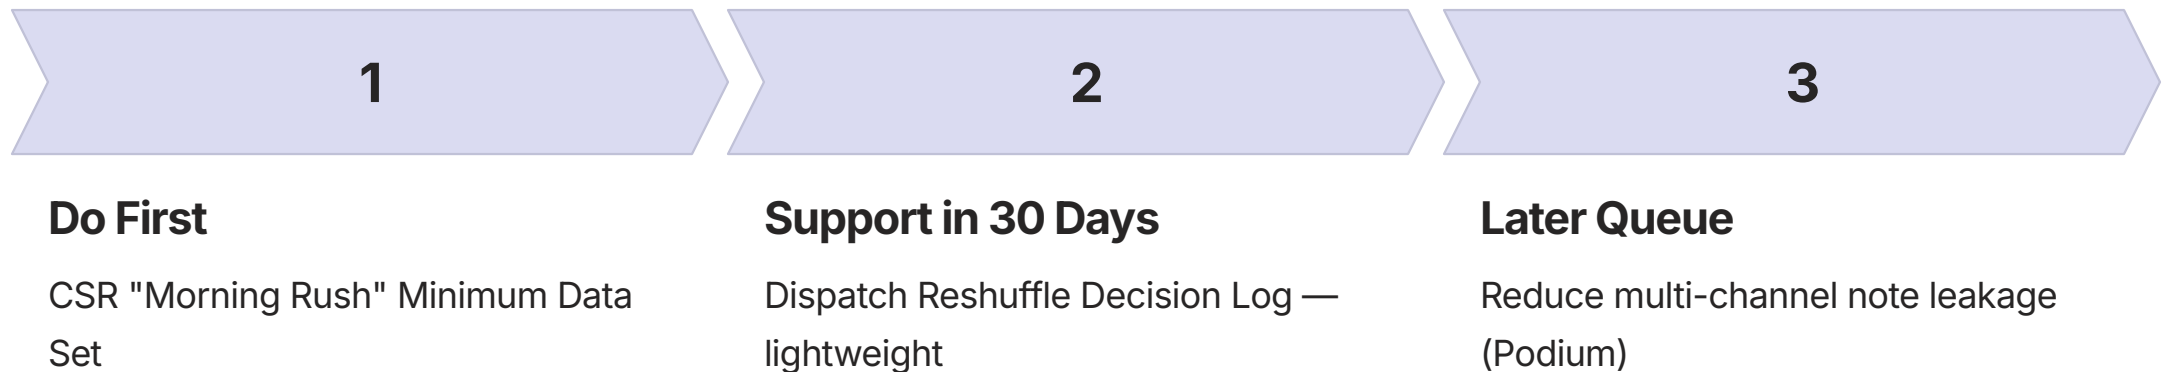
## Risks & Governance

Recommendation risks, governance notes, and appendix steps live in the web report.

 **Evidence Note:** Summary only. Use the web report or Markdown source for full evidence, risks, and implementation detail.

# Executive Answer

Run a 10–14 day workflow readiness + implementation push focused on 1 CSR minimum data set.



## Wait On — Do Not Proceed Yet

- Fully autonomous AI dispatching or schedule reshuffling — high risk
- AI-driven technician performance scoring or automated coaching recommendations based on current callback data
- Automating Podium/CallRail ingestion into AI summaries
- Replacing ServiceTitan as the system of record — explicit constraint

**Support option:** Build.

# Current Workflow Snapshot

Keep the deck tied to workflow evidence, not generic AI ideas.

01

## Trigger

Inbound call or message received.

03

## Urgent Escalation

If urgent (no heat / no cooling), CSR messages dispatch directly.

05

## Assign in ServiceTitan

Dispatcher assigns job in ServiceTitan.

02

## CSR Answers

CSR captures address, symptoms, safety notes, and availability.

04

## Dispatcher Reshuffles

Dispatcher reshuffles schedule — mostly in head — considers tech strengths.


06

## Tech On-Site → Close Out

Tech arrives on-site and closes out job. If issue persists, callback logged (may be new).

## Tools in Use

ServiceTitan · QuickBooks

 **Evidence Limit:** Which ServiceTitan edition/modules Summit has — templates, required fields, custom fields, internal notes — is not confirmed.

# Workflow Snapshot — Evidence Limits

The following data points were not available at the time of this audit. These gaps should be resolved before expanding automation scope.

## Ownership

Who is the named owner for each workflow step.

## Baseline Volumes

Call volume, dispatch volume, and callback frequency baselines.

## ServiceTitan Fields

Current ServiceTitan fields/templates used by CSRs.

## CallRail Recording

Whether CallRail calls are recorded.

## Podium Export

Whether Podium messages/photos can be exported via API.

## Callback Reason Codes

Current list of callback reason codes, if any exist.

- How membership priority rules work operationally — what exactly changes
- How tech "strengths" are represented today — tags, service lines, certifications
- Stakeholder map for approvals
- Data governance requirements

# Pain Points & Constraints

☐ These are the limits that should shape the first pilot.

## Fragile Dispatch Operations

**Pain:** Single-dispatcher tacit knowledge — reshuffle logic lives in one person's head.

**Risk:** Fragile operations when dispatcher is overloaded or out.

## CSR Intake Inconsistency

**Pain:** Inconsistent intake during peak demand.

**Risk:** Downstream mis-prioritization of jobs.

## Multi-Channel Comms

**Pain:** Text/Podium messages alongside calls.

**Risk:** Higher chance of incomplete ServiceTitan notes.

## Callback Data Quality

**Pain:** Callback reporting data quality is low.

**Risk:** Leaders distrust performance insights.

## Change Resistance

**Pain:** Resistance to "more admin."

**Risk:** Anything that adds steps outside ServiceTitan or adds call time is likely to be rejected.

⚠ **Evidence Note:** E1, E2, E3, E7, E11

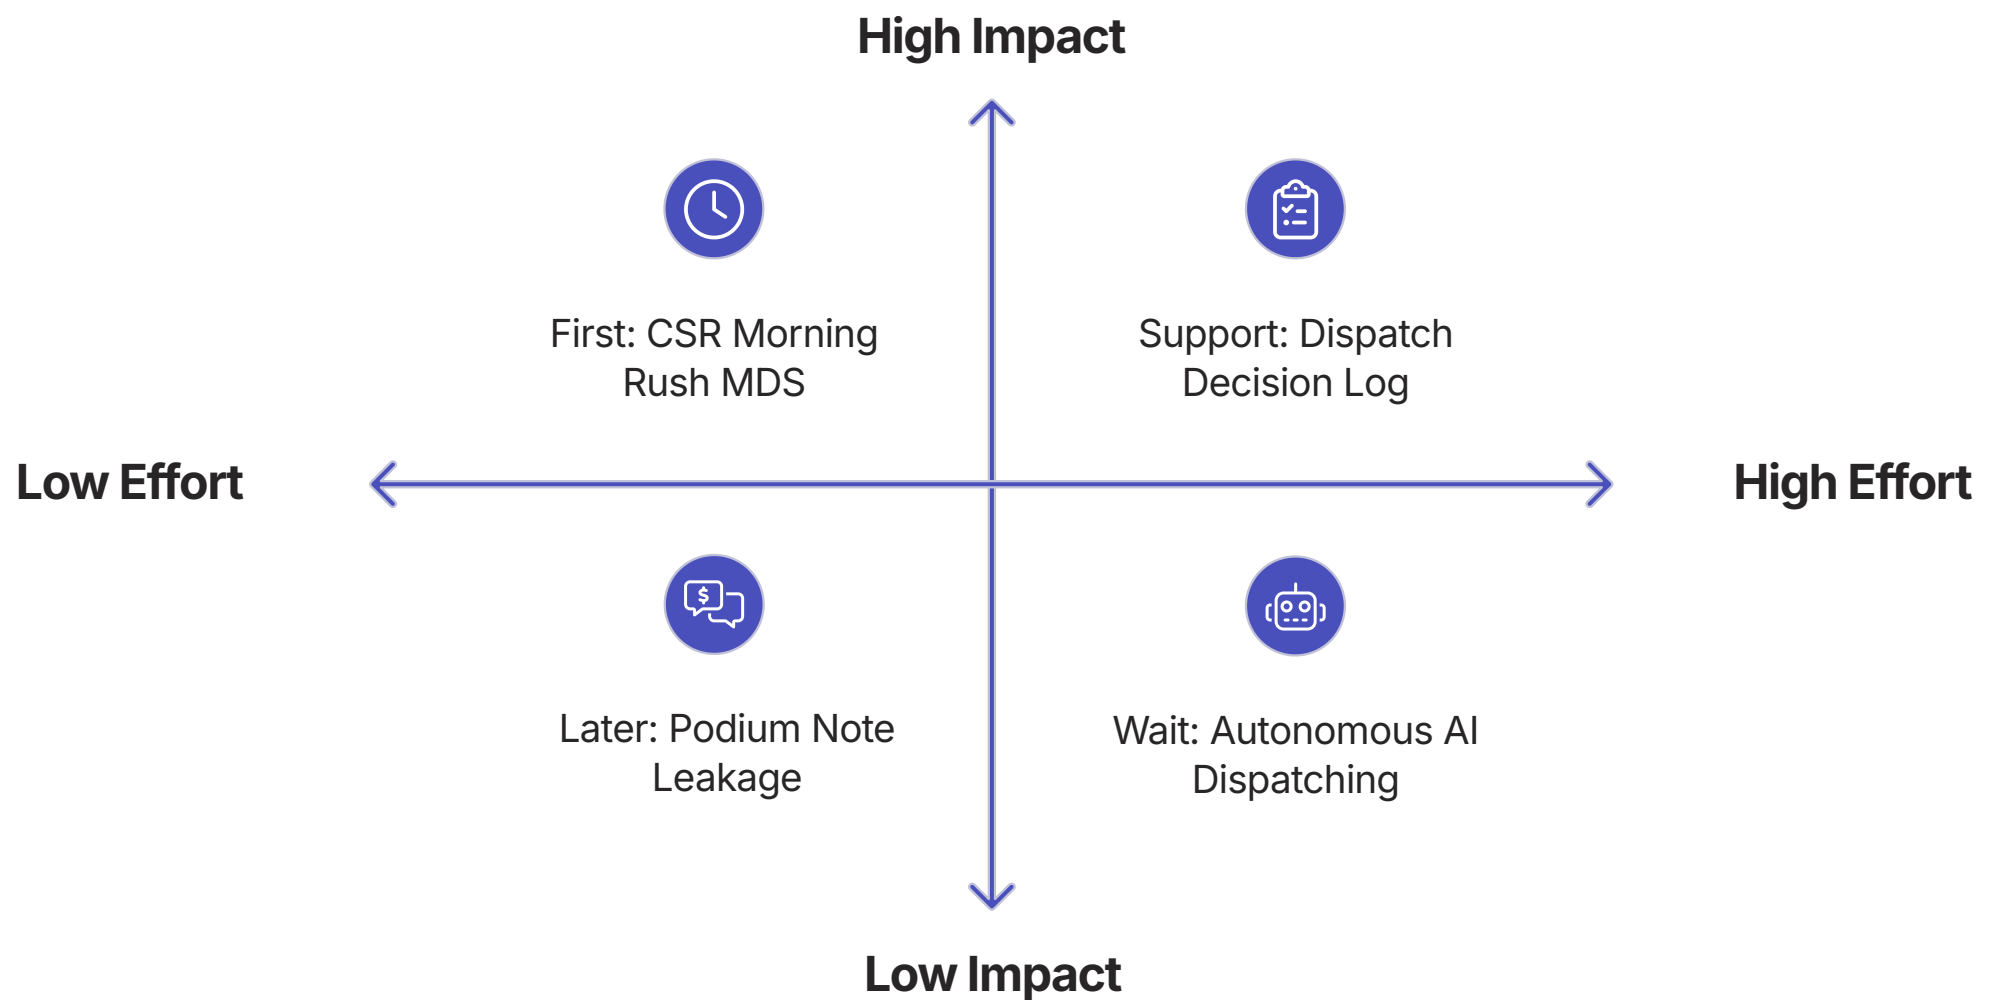
# Priority Matrix

Ordered for first pilot, 30-day support, and later queue.

Queue	Move	Impact	Confidence	Payoff
First	CSR "Morning Rush" Minimum Data Set	High	Medium	Improved intake completeness for membership/equipment context.
Support	Dispatch Reshuffle Decision Log — lightweight	High	Medium	Reduced operational fragility.

# Priority Matrix — Continued

Queue	Move	Impact	Confidence	Payoff
Support	Callback Taxonomy	Medium	Medium	Trusted, apples-to-apples callback reporting that can support improvement work.
Later	Reduce multi-channel note leakage (Podium)	Medium	Medium	Fewer missed commitments.



The priority sequence moves from highest-confidence, lowest-risk actions first — building data quality before layering on automation.

FIRST MOVE

# CSR "Morning Rush" Minimum Data Set

Pilot payoff is improved intake completeness for membership/equipment context.

## Pilot Action

Pilot CSR "Morning Rush" Minimum Data Set for 30 days.

## Owner

General Manager Maria Chen + CSR Lead (not named) +  
Dispatch Lead

## Start Condition

GM + dispatch lead agree on MDS fields before launch.

## Why This First?

Automated summaries or flags do not help if core fields aren't reliably captured. Data quality must come before automation.

## Technology Path

Existing stack configuration — the required information already exists in ServiceTitan.

# CSR "Morning Rush" MDS — Measurement & Control

## Measure

Weekly MDS completeness rate on urgent calls + weekly count of dispatch/tech escalations.

## Control

Human review gate — no automated enforcement during pilot.

## Tools Now


ServiceTitan configuration only.

## Support Option

Build.

## Evaluate Later

After MDS capture is stable, evaluate automation platform to draft summaries from Podium.

 **Why Not More Automation Now:** Automated summaries or flags do not help if core fields aren't reliably captured first.

# CSR "Morning Rush" MDS — Evidence Note

Specific downstream-critical fields — membership status, equipment age/replacement signal — are frequently missed in notes.

## The Core Problem

When CSRs miss membership or equipment age fields during intake, dispatchers and technicians arrive on-site without the context needed to prioritize or prepare correctly.

## Why It Matters Downstream

Missed intake fields create mis-prioritization, longer on-site time, and callback risk — all of which compound during peak demand periods like morning rush.

## Why MDS Is The Right Fix

A minimum data set enforced at the ServiceTitan booking template level captures the right fields without adding significant call time — addressing change resistance directly.

# Dispatch Reshuffle Decision Log

Operational payoff is reduced fragility.

## Pilot Action

Pilot Dispatch Reshuffle Decision Log — lightweight — for 30 days.

## Owner

Dispatch Lead (senior dispatcher, not named) + GM Maria Chen

## Start Condition

Agree on threshold for what constitutes a "significant reshuffle" before logging begins.

## The Fragility Problem

Reshuffle logic lives entirely in one dispatcher's head. When that person is overloaded or absent, operations become fragile.

## The Log Solution

A lightweight decision log begins to externalize tacit knowledge without requiring full process redesign.

# Dispatch Decision Log — Measurement & Control

## Measure

% of significant reshuffles that have a completed log entry.

## Control

Keep the log operational — no HR/performance language.  
Require human dispatcher approval on all entries.

## Technology Path

Workflow readiness first — the bottleneck is undocumented decision logic, not missing technology.

## Support Option

Enable.

---

## Tools Now

Workflow readiness first — no new tooling required to start.

## Evaluate Later

After the log is consistently used, evaluate whether pattern analysis or AI-assisted suggestions are appropriate.

# Dispatch Decision Log — Why Not More Automation

Automated dispatch recommendations are high-risk. Without a human-readable decision trail, Summit cannot validate, audit, or safely expand AI involvement.

## Single Point of Failure

Dispatch knowledge is currently concentrated in one person. This is the primary operational risk identified in the audit.

## Log Before Automate

The decision log creates the paper trail that makes future automation safe. Skipping this step and going straight to AI dispatch recommendations would remove human oversight before the logic is even documented.

## Human Approval Required

All reshuffle decisions must retain human dispatcher approval throughout the pilot period.

 **Evidence Note:** Single-point-of-failure dispatch knowledge.

SUPPORT MOVE

# Callback Taxonomy

Primary payoff is trusted, apples-to-apples callback reporting that can support improvement work.

## Pilot Action

Pilot Callback Taxonomy for 30 days.

## Owner

GM Maria Chen + Service/Ops  
Manager (not named) + Dispatch  
Lead

## Start

Identify the manager group that owns  
definitions before any coding begins.

## Measure

Sample audit of callback entries  
against agreed reason codes.

## Control

Written governance statement +  
manager review requirement.

## Evidence

Attribution ambiguity + trust gap in  
current callback reporting.

# Reduce Multi-Channel Note Leakage (Podium)

Pilot payoff is fewer missed commitments.

## Pilot Action

Pilot Reduce multi-channel note leakage (Podium) for 30 days.

## Owner

CSR Lead + Dispatch Lead

## Start

Add the comms summary block to the urgent call template in ServiceTitan.

## Measure

% of urgent jobs with a completed Customer Comms Summary in the job note.

## Control

Summarize — do not paste raw message threads.

## Why This Matters

Multi-channel customer communication (calls + texts + Podium) increases cognitive load on CSRs and dispatchers, raising the chance that commitments made in one channel are never recorded in ServiceTitan.

## Privacy Note

Pasting raw message threads into job notes may create privacy concerns. Summarize only.

# 30-Day Quick Wins

✔ These are usable actions, not promises for future consulting.

## 1 Deploy the CSR 30-Second Script

Use verbatim for urgent calls to reliably capture the minimum data set fields every time.

## 2 Add a 'MissingInfo' Line

Add to the ServiceTitan urgent-call note block so gaps are visible to dispatchers immediately.

## 3 Run End-of-Day Dispatch Review

10-minute review for 5 reshuffles to confirm decision log entries are accurate and complete.

## 4 Publish Callback Reason-Code List v1

A short, agreed list of reason codes that all managers use consistently.

## 5 Adopt Three-Line 'Customer Comms Summary'

Inside ServiceTitan for any job where customer communication occurred across multiple channels.

# Starter Artifacts

The following artifacts are referenced in this deck. Full versions remain in the report appendix.

## Fields-to-Track List

Morning Rush CSR Intake — Minimum Data Set (MDS) Goal.  
Defines which fields must be captured on every urgent call.

## Decision Log Template

Dispatch Reshuffle Decision Log — v1 copy/paste. Includes  
"when to use" v1 rule for defining a significant reshuffle.

## Callback Scorecard

Callback Definitions + Attribution Rubric — v1 one-page.  
Includes callback definition and attribution guidance for  
managers.

## Note Template

ServiceTitan Job Note Insert — Customer Comms  
Summary copy/paste. Standardizes the three-line summary  
format for multi-channel jobs.

 **Evidence Note:** Full starter artifacts remain in the report appendix.

# First 30 Days

- Keep the pilot narrow enough to measure and adjust.

## CSR Morning Rush MDS

ServiceTitan 'Urgent Call' booking template live with MDS fields configured.

1

## Callback Taxonomy

Callback rubric agreed by managers. Written governance statement in place.

2

3

4

## Dispatch Decision Log

Decision log template in use for urgent insertions. Human dispatcher approval required.

## Later Queue — Out of Scope

Reduce multi-channel note leakage (Podium) + calls by standardizing a 'single source' job-note handoff kept out of scope for first 30 days.

# What Comes Later

## Days 31–60

- Review CSR "Morning Rush" MDS + required fields in ServiceTitan booking using weekly completeness data.
- Review Dispatch Reshuffle Decision Log to reduce single-point-of-failure using % of significant reshuffles with decision entries.
- Review Callback Taxonomy + attribution rules so leaders can trust the data and avoid "AI policing."
- Revisit the later queue only after the first pilot has reliable data and a clear owner.

## Days 61–90

- Expand CSR "Morning Rush" MDS + required fields in ServiceTitan booking only if pilot data supports it.
- Expand Dispatch Reshuffle Decision Log only if usage, review quality, and source data are stable.
- Expand Callback Taxonomy + attribution rules so leaders can trust the data.
- Consider sequencing Reduce multi-channel note leakage (Podium) + calls by standardizing a 'single source' job-note handoff.

# Evidence Behind The Deck

The following evidence items underpin the recommendations in this deck. Full evidence detail is in the web report.

## E1

Single-point-of-failure dispatch knowledge.

## E2

CSR intake inconsistency.

## E3

Multi-channel customer communication increases cognitive load.

## E4

Dispatching depends on tacit skill-matching heuristics.

## E5

Specific downstream-critical fields — membership, equipment age/replacement signal — are frequently missed in notes.

## E6

Business impact of missed intake fields.

# Risks To Control

## Adoption Risk

If required fields add noticeable call time, CSRs may resist or skip them. Time-box the note step explicitly.

## Dispatch Skip Risk

If the note step is not time-boxed, dispatch may skip it during high-volume periods.

## Definition Dispute Risk

If definitions are not agreed across managers, callback data will remain disputed and unusable.

## Privacy Risk

If CSRs paste raw message threads into job notes, it may create privacy concerns. Summarize only.

---

## Open Evidence Gaps That Increase Risk

- Which ServiceTitan edition/modules Summit has — templates, required fields, custom fields, internal notes
- Who is the named owner for each workflow
- Baseline volumes
- Current ServiceTitan fields/templates used by CSRs
- Whether CallRail calls are recorded
- Whether Podium messages/photos can be exported via API
- Current list of callback reason codes if any
- How membership priority rules work operationally
- How tech "strengths" are represented today — tags, service lines, certifications
- Stakeholder map for approvals

# More Risks To Control

- ⊗ **Open Evidence Gap:** Data governance requirements for Summit Comfort HVAC have not been confirmed. This gap must be resolved before expanding any automation scope beyond the first pilot.

Until data governance requirements are documented and agreed, any automation that touches customer data, technician performance data, or call recordings carries unquantified compliance risk.

## What This Means for the Pilot

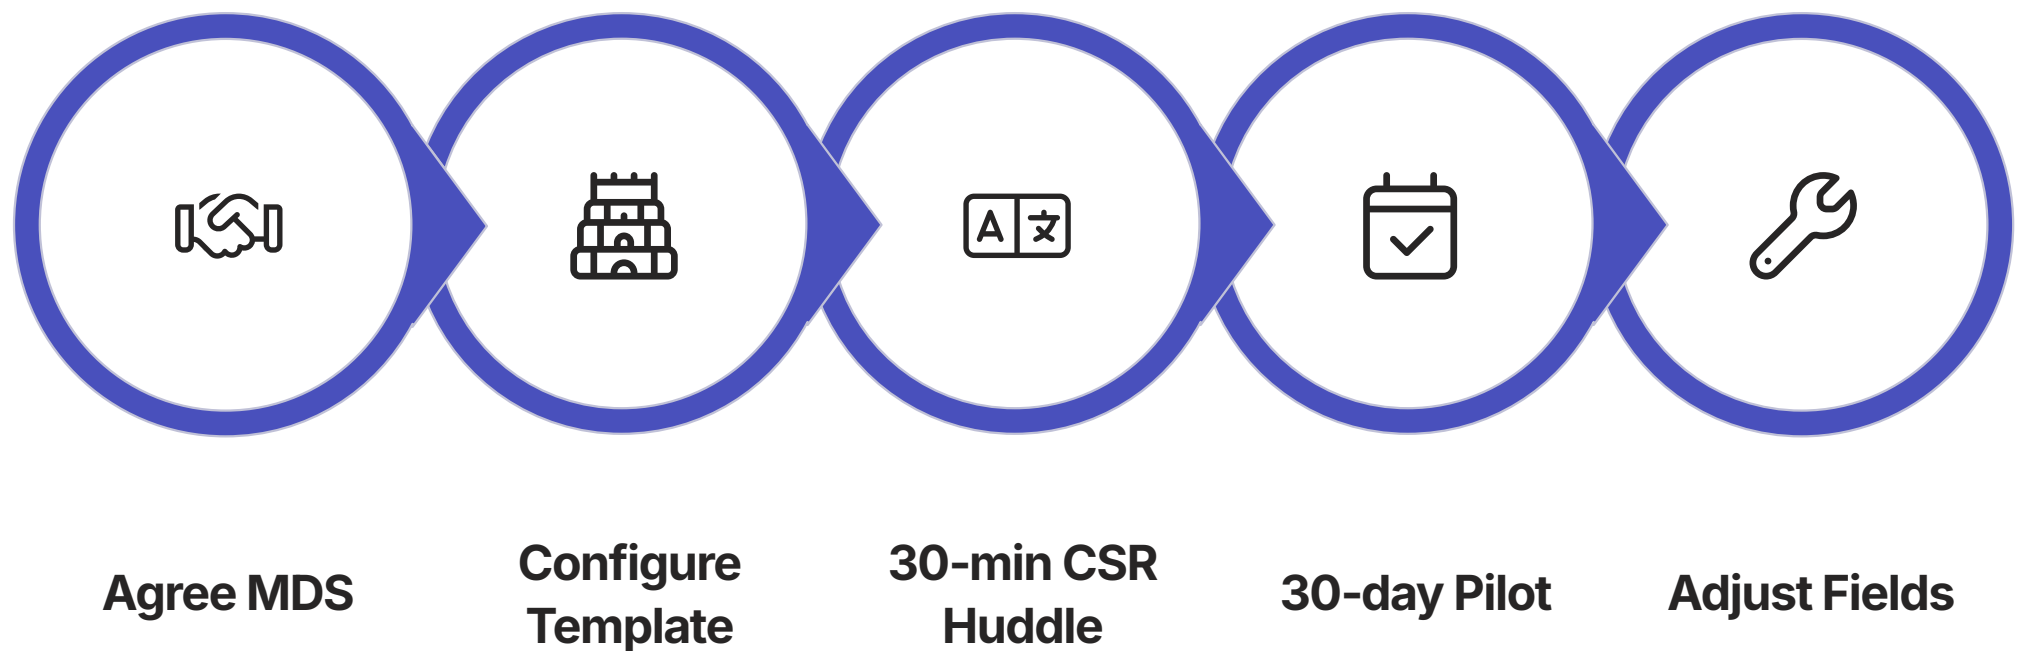
The first pilot (CSR MDS + Dispatch Decision Log) operates within ServiceTitan's existing data model and does not require new data governance decisions. This is one reason it is the recommended first move.

## What This Means for Later Queue Items

Automating Podium/CallRail ingestion, AI-driven performance scoring, and multi-channel summary generation all require data governance clarity before proceeding.

# Recommended Next Step

Run a 10–14 day workflow readiness + implementation push focused on 1 CSR minimum data set.



This sequence keeps the pilot narrow, measurable, and adjustable — matching the change resistance constraint identified in the audit.

## Next Action

Pilot CSR "Morning Rush" Minimum Data Set for 30 days.

## Named Owner

GM Maria Chen + Dispatch Lead

## Supporting Evidence

Named owner, measurable version, and supporting evidence are all in place. Start here.

**Evidence Note:** Start with CSR "Morning Rush" Minimum Data Set: named owner, measurable version, supporting evidence confirmed.