



AI READINESS ASSESSMENT

A ranked plan for where AI belongs, where it does not, and what to prove first.

Worked specimen for a fictional 25-person field services company.

ONE-WEEK DIAGNOSTIC

What this specimen demonstrates

- A defensible workflow inventory
- Priorities with reasons and stop conditions
- A bounded pilot brief and 90-day sequence
- Data handling and operational controls

Important: every company name, workflow, volume, score, estimate, and recommendation in this document is fictional. It shows the format and standard of reasoning, not a client result or a promised outcome.

Prepared 2026-07-25

Specimen version 1.0

Decision artifact, not a sales estimate

The real assessment is based on interviews, observed workflow evidence, documented constraints, and the client's own acceptance thresholds.

DECISION IN ONE PAGE

Start with intake triage. Defer autonomous customer replies.

This fictional company has enough process consistency to pilot assistive automation, but not enough control evidence to let a model send anything without review.

ILLUSTRATIVE READINESS

64 / 100

Ready for one bounded pilot

TOP CANDIDATE

#1 Intake

High repetition, reviewable output

PRIMARY CONSTRAINT

Data rules

Customer PII needs a written boundary

Recommended first move

- Pilot extraction and routing for inbound service requests.
- Keep a human approval step before records are created or messages are sent.
- Measure review time and error categories across 50 representative requests.
- Stop if the team cannot agree on one owner or one source-of-truth queue.

What not to buy yet

- A general-purpose chatbot for customers
- A new CRM before current fields are cleaned up
- An enterprise model platform before a pilot proves demand
- Automated dispatch decisions with no appeal path

Readiness score - illustrative calculation

Dimension	Weight	Score	Evidence expected in a real assessment
Process repeatability	30%	72	Observed steps, stable inputs, named exceptions
Data controls	25%	48	Approved systems, retention, access and vendor rules
Ownership	20%	70	One accountable owner and escalation path
Measurement	15%	58	Baseline volume, review time, error taxonomy
Change capacity	10%	76	Available pilot users and training window

WORKFLOW EVIDENCE

Observed work, normalized before recommendations

The assessment ranks recurring work only after documenting volume, variation, data sensitivity, and the cost of a wrong answer.

#	Workflow	Illustrative evidence	AI fit	Decision
1	Service request intake Email to dispatch queue	85/week; 6 common fields; 12 exception types; customer PII	High	Pilot assistive extraction and routing
2	Weekly operations report Three systems to one deck	4.5 hr/week; stable template; numbers must reconcile	High	Automate collection; keep variance review
3	Invoice follow-up Aging list to outreach	45/week; financial data; message tone varies by account	Medium	Draft only after data cleanup
4	Knowledge search SOP and policy lookup	18 docs; conflicting versions; no named content owner	Medium	Fix ownership and archive stale files first
5	Customer replies Status and scheduling questions	High consequence if wrong; account context in two systems	Low now	Defer autonomous replies
6	Dispatch decisions Assign crew and time window	Safety, skills and route constraints; frequent exceptions	Low	Rules and dashboard first, no model decision

How to read the ranking

Evidence before enthusiasm

A workflow is not a candidate because it sounds repetitive. The report names the volume, stable inputs, exceptions, and cost of failure.

Assist before automate

A reviewable draft or extraction step earns trust faster than an end-to-end autonomous action.


A deferral is a deliverable

Items 5 and 6 are explicit no-go recommendations until the listed controls exist.

DECISION LOGIC

Payoff alone does not determine priority

The first pilot sits where payoff, feasibility, reversibility, and evidence quality overlap. Bubble size represents operational risk, not project price.



Ranked sequence

- 1 Intake triage
- 2 Weekly report
- 3 Invoice follow-up
- 4 Knowledge search
- 5 Customer replies
- 6 Dispatch decisions

Why #1 wins

Its output is structured, easy to compare with a human answer, and reversible before any downstream action. It creates evidence the next decision can use.

Red bubbles are not hidden opportunities. They are attractive outcomes with high consequence and weak current controls. The recommendation is to improve the process around them, not to automate them first.

BOUNDED EXPERIMENT

Pilot 01 - service request intake assistant

A two-week, human-in-the-loop proof that extracts six fields, suggests a route, and explains uncertainty without creating records or contacting customers.

Current state

- Shared inbox receives free-form requests.
- Coordinator reads, copies six fields, and chooses a queue.
- Missing details trigger a manual follow-up.
- Exceptions are remembered by individuals, not logged.

Proposed state

- Parse a copy of the inbound message.
- Redact or isolate restricted fields before model use.
- Return six fields, route suggestion, confidence, and reason.
- Coordinator approves or corrects; corrections become test evidence.

Acceptance evidence - agreed before build

Measure	Illustrative threshold	How it is checked
Field extraction	At least 95% across 50 labeled requests	Exact-match score by field and exception type
Routing suggestion	At least 90%; never auto-applied	Compare suggestion with coordinator decision
Review time	Median at or below 2 minutes	Timestamp from suggestion to approve/correct
Unauthorized actions	Zero	No write credentials and no outbound-send path
Explainability	Reason and confidence present on every case	Schema validation before result is shown

Stop conditions

Stop and redesign if test data cannot be legally approved, if exception labels are not reproducible between reviewers, or if any proposed architecture requires direct write access to dispatch during the proof.

RISK BOUNDARY

Data boundary before tool selection

The report separates what may enter a cloud model, what needs a contractual decision, and what stays out. The real version names approved vendors and owners.

Data class	Example	Default handling	Required control
Public	Published service descriptions	Cloud tools allowed	Source citation and human review
Internal	SOPs, queue names, internal metrics	Approved tenant only	No training; retention configured; access logged
Customer PII	Name, address, phone, service history	Redact or isolate by default	Written purpose, approved processor, deletion path
Financial	Invoice status and balance	No model access in first pilot	Least privilege and reconciliation evidence
Credentials	Passwords, API keys, access tokens	Never place in prompts	Secret manager; rotation; scoped service identity

Minimum control set for the pilot

1

No production writes

Read-only copy or synthetic fixture data.

2

Human approval

Every suggestion is accepted or corrected before action.

3

Traceable output

Log input ID, model/version, schema result, reviewer decision.

4

Bounded retention

Define what is stored, for how long, and who can delete it.

5

Failure path

Timeouts and low confidence return work to the existing queue.

6

Named owner

One operator owns labels, exceptions, and escalation.

Approval gate

The pilot does not move to live customer data until the data owner approves the processor, retention, access, deletion, and incident path in writing.

SEQUENCED DELIVERY

Prove one thing, then earn the next decision

Each phase ends with evidence and an explicit stop or continue choice. Dates in a real report are anchored to kickoff and named owners.

WEEKS 1-2

Bound

Name owner, approve data, label 50 cases, document exceptions.

Evidence: approved test set and baseline review time.

WEEKS 3-4

Build

Implement read-only extraction, schema checks, review screen, and logs.

Evidence: end-to-end test in an isolated environment.

WEEKS 5-8

Prove

Run shadow mode, compare corrections, and review weekly error categories.

Evidence: threshold table with pass or fail by field.

WEEKS 9-13

Decide

Adopt, redesign, or stop. Only then scope reporting automation.

Evidence: signed decision record and handover package.

Portfolio rule

Do not start a second AI pilot while the first has unresolved data ownership, measurement, or adoption failures. Parallel experiments create activity, not evidence.

REUSABLE METHOD

The report is designed to be executable without Parker AI

A buyer owns the inventory, assumptions, priority logic, pilot brief, and controls. Implementation with Parker AI is optional.

01

Discover

Walk the work with the people who perform it.

02

Bound

Name data, consequence, owner, and exceptions.

03

Rank

Score payoff, feasibility, reversibility, and evidence.

04

Prove

Define thresholds and stop conditions before building.

05

Hand over

Deliver the plan, decision record, and next evidence.

Assumption ledger

Assumption	How to verify
85 requests/week	Export 8 representative weeks
Six stable fields	Compare form and queue schema
2-minute target	Time current review sample
One queue owner	Name owner and backup in kickoff

Handoff package

- Executive readout and workflow inventory
- Ranked opportunity map with deferrals
- Pilot brief, acceptance thresholds, and stop conditions
- Data map and control checklist
- 90-day sequence and owner decisions
- Editable source notes and assumptions

This document is a composite example, not a real client report.

Its numbers and recommendations are invented solely to demonstrate structure. A real assessment may conclude that AI is not the right next investment.

getparkerai.com