

AI Implementation Playbook for Service Businesses

Version: 0.1.0

Format: Working source draft for PDF/long-read production

Downloadable diagrams

- `diagrams/content-workflow-diagram.svg`
- `diagrams/support-workflow-diagram.svg`
- `diagrams/lead-routing-workflow-diagram.svg`

Interview validation assets

- `ai-implementation-operator-interview-guide.md`
- `ai-implementation-interview-synthesis-template.md`

1) Introduction and safety baseline

What this playbook is and is not

- Operator-focused implementation guide for service businesses.
- Not a benchmark sheet or static "best model" ranking.
- Designed for measurable workflow gains in 30 days or less.

Safety and data hygiene

- Classify data before tool trials (public/internal/restricted/regulated).
- Define restricted prompt/export fields before production use.
- Require human review for high-impact outbound communication.

Governance

- Who approves: Ops lead plus legal/compliance when data sensitivity requires it.
- What to log: Data classes, restricted fields, mandatory review checkpoints.

Case insert (before -> after)

- Before: AI support drafting began with no redaction rules.
- After: Team enforced data classes and reviewer gates; no sensitive leakage in outbound drafts.

2) Assess your stack

System map

- Map where customer interactions start/end across inbox, CRM, calendar, docs, and phone.
- Document handoffs and latency between humans and systems.
- Identify duplicate capture points creating rework.

Friction inventory

- List recurring operational failures (response lag, poor qualification, lost notes).
- Estimate weekly volume and business cost for each failure mode.
- Prioritize high-volume, low-exception workflows first.

Governance

- Who approves: Ops manager with sales/support/content owners.
- What to log: Current-state map, baseline metrics, known system constraints.

Case insert (before -> after)

- Before: Qualification happened in multiple systems with conflicting criteria.
- After: Unified CRM rules reduced reassignment and first-response variance.

3) Use case selection and scoring

Scoring model

- Score use cases on volume, error cost, data availability, and change burden.
- Choose one workflow pilot at a time for clean measurement.
- Avoid bundled launches that obscure causality.

Pilot definition

- Run 2-4 week pilots with success and rollback conditions defined upfront.
- Document expected failure modes before launch.
- Assign one accountable owner for go/no-go decisions.

Governance

- Who approves: Founder or head of operations.
- What to log: Scoring sheet, hypothesis, success metrics, rollback triggers.

Case insert (before -> after)

- Before: Three workflows launched together with unclear outcome attribution.
- After: Sequenced pilots tied each metric change to one intervention.

4) Functional playbooks

Content

- Use AI for first drafts and repurposing; keep a fixed quality rubric.
- Require human review for final claims and brand tone.
- Capture revision reasons to improve prompts.

Support

- Use confidence-gated triage labels and macro suggestions.
- Escalate low-confidence/high-impact cases automatically.
- Track escalation rates and SLA impact weekly.

Lead routing

- Define qualification signals and response-time SLAs.
- Log handoff ownership and timestamps in CRM.
- Alert humans when routing confidence is low.

Forecasting and operations

- Use range-based scenarios; avoid false precision.
- Separate leading indicators from lagging revenue outcomes.
- Review and refresh assumptions monthly.

Governance

- Who approves: Functional leader for each workflow plus shared ops owner.
- What to log: Prompt changes, escalation volume, SLA trends, assumption updates.

Case insert (before -> after)

- Before: Support auto-replies were sent without confidence thresholds.
- After: Human gating reduced incorrect responses and stabilized CSAT.

5) Rollout plan and vendor decisions

Rollout cadence

- Pilot -> measure -> expand.
- Publish weekly operator notes (what changed, what broke, next fix).

- Expand only when training and ownership are clear.

Vendor/model criteria

- Evaluate reliability, governance controls, integration fit, and operating cost.
- Prefer criteria-based decisions to avoid stale "best model" claims.
- Keep an exit path to reduce lock-in risk.

Governance

- Who approves: Founders and operations leadership with finance review.
- What to log: Selection criteria, trade-offs, incident notes, retraining updates.

Case insert (before -> after)

- Before: Tool selected from demo quality alone; integration failed post-launch.
- After: Criteria-first selection improved inbox/CRM reliability.

Appendix: readiness checklist

- Name highest-volume repeat work across content/support/routing/ops.
- Assign accountable human owner for each workflow.
- Confirm access to source systems and sample data.
- Define low-confidence escalation paths.
- Define launch metrics before pilot start.

Further reading prompts

- Which process has the highest weekly manual volume?
- Where do quality failures create downstream cost?
- Which handoff delay most impacts conversion or retention?