

Product leadership that connects strategy, operations and AI

Omnichannel retail • Aviation • Payments • B2B SaaS

I turn ambiguous customer and operational problems into measurable launches, then build the evidence to scale them.

Sydney • bhaskar.j.ashoka@gmail.com • [linkedin.com/in/bhaskarashoka](https://www.linkedin.com/in/bhaskarashoka)

The work in one view

Direct to Boot

681 → 755 sites across F22–F25
~8 → 4 minutes median handover
Growth roadmap, launch gates
and service quality

Qantas servicing

~45% front end latency reduction
24 high risk scenarios mapped
Native cancellation and refund
journey

Payments and SaaS

34% faster incident detection at
IR
50+ enterprise accounts
Commercial cards and B2B
payment flows

Figures on these slides are drawn from the latest supplied resume; opportunities and prototypes are labelled separately.

Direct to Boot: scale with a service gate

Problem	Rapid pickup expansion made store readiness and arrival signalling critical.
My role	Owned growth roadmap, adoption measures, site readiness and cross team tradeoffs.
Decision	Use demand, bay capacity, training, arrival events and exceptions to launch, hold or remediate.
Outcome	681 sites in F22 to 755 in F25; ~8 to 4 minute median handover.

The site growth trend is not attributed solely to one intervention.

A commercial model that survives scrutiny

Capacity opportunity

4 minutes saved × eligible collections = theoretical service minutes released.

Convert to paid labour only if staffing schedules, observed volumes and Finance evidence support it.

Scale guardrails

Median and p90 handover, completed collections, manual fallback, errors, complaints, staff workload and site readiness.

Compare pilot and control stores before claiming incremental impact.

Example: 180 collections/day × 4 minutes = 720 minutes/day of theoretical capacity; not a realised payroll saving.

Retention: targeted Rewards experiment

Experiment design

Segment eligible pickup members. Test targeted bonus points and re engagement against a control group. Monitor offer cost, service quality and repeat behaviour.

Reported result

90 day repeat purchase: 31.0% → 34.7% in targeted cohorts.
+3.7 percentage points, approximately +12% relative.
Avoid extrapolating to all customers.

Decision rule: expand only if incremental contribution and fulfilment quality justify the incentive.

Qantas: native cancellation and refunds

Product decisions

Map customer eligibility, card refunds, vouchers, partial use, mixed payments and disrupted itineraries. Convert 24 high risk scenarios into API requirements and operational fallbacks.

Result and opportunity

Approximate p95 front end latency: 5.5 → 3.0 seconds (~45%).

Estimated 35% eligible contact deflection was the business case opportunity, not a measured outcome.

Commercial test: eligible contact volume × deflection rate × cost/contact, adjusted for channel shift and exceptions.

Why an agent workflow here?

The useful output is a reviewable decision: evidence, recommendation, owner, guardrails and a measurement plan.

Observe

Collect permissioned events, validate freshness and detect an exception.

Reason

Compare options, quantify uncertainty and apply fixed policies.

Act and learn

Require approval, log the action and compare actual outcomes.

The linked demos run on synthetic data and deterministic rules. Production LLM and tool integrations are proposed architecture.

Prototype A: Retail Operations Copilot

End to end path

- 1 Detect median handover above target
- 2 Diagnose arrival signalling or staging risk
- 3 Estimate theoretical capacity released
- 4 Propose a two week store pilot
- 5 Ask store lead to approve
- 6 Compare pilot with control stores

Approval and safeguards

Keep a manual arrival path. Review location consent. Monitor p90 wait, errors, complaints and staff workload. Store lead owns the decision. No external action occurs in the browser demo.

Prototype B: Payments Migration Copilot


End to end path

- 1 Capture billing, volume and reconciliation needs
- 2 Map authorisation, settlement and refunds
- 3 Flag stored credential and mandate risks
- 4 Build migration backlog and sandbox tests
- 5 Ask risk, finance and engineering to approve
- 6 Pilot cohorts with rollback

Measures and safeguards

Watch authorisation, disputes, churn, reconciliation breaks, cost per successful payment and incidents. Never infer savings without fee and mix data. No payment details are entered or transmitted.

A production architecture with bounded actions



Start read only. Use server side secrets, least privilege, data minimisation, redaction and source citations. Human approval precedes any write operation.

Production implementation requires security, privacy, provider and operational review.

How I would build and validate the real product

Days 1–30 · Discover

Select one workflow. Audit data access and policy boundaries. Establish a manual baseline and 20–30 evaluation cases.

Days 31–60 · Pilot

Build read only connectors, orchestration and approval UI. Test accuracy, uncertainty, false positives, latency and cost.

Days 61–90 · Decide

Run controlled pilot. Compare time to decision, quality, incidents and contribution. Expand only when guardrails hold.

Exit criteria are agreed with operational, security and finance owners before a live pilot.

Let's build something useful

I bring product judgement, technical fluency and the operating discipline to turn a prototype into a measured service.

Bhaskar Ashoka · Sydney, Australia

bhaskar.j.ashoka@gmail.com · linkedin.com/in/bhaskarashoka