proservartner

# Building & Orchestrating *AI Agents* in GBS.

Most GBS programmes don't fail at the agent. They fail at the **sequence** - picking the wrong first use case, scaling before a platform exists, or running pilots without a strategy that survives them. This guide sets out ***the seven steps we walk clients through***, from strategy and operating model into operating rhythm, plus the cross-cutting issues that quietly stall programmes.

WHAT'S INSIDE


01 Strategy & *operating model*
02 Discovery via *task mining*
03 The *use case library*
04 Pilot *quick wins*
05 Scale *end-to-end*
06 Industrialise *the platform*
07 Embed & *evolve*
▸▸ 12 use cases *to start with*


Built for CFOs, GBS leaders and shared-service directors. Worked through with finance, but the same sequence holds for HR, procurement, supply chain and IT. Typical engagement: ***12 to 18 months*** from sponsor sign-off to second cycle in production.

FOUNDATION

# Steps 1-4 - *before any agent goes live.*

## 01 Strategy & *operating model*

Ten dimensions we walk clients through to **decide what GBS becomes** - before any tooling.

- **Inspire the need** - the case for change in GBS and across the wider organisation
- **Sourcing** - captive, BPO, hybrid; the right sourcing call
- **Processes & organisation** - end-to-end, global ownership, waste out
- **Guardrails** - misuse, privacy, fairness, accountability, compliance
- **HL business case** - returns, costs, funding ask, sponsor commitment
- **AI governance framework** - model risk, audit trails, decision rights
- **Culture & training** - embrace AI; certification and uplift framework
- **Agentic AI & automation** - right tools, right way, high-value implementations
- **Customer orientation** - internal customer at the front of every build
- **Competency & talent** - reshape jobs, new skills, attract talent

KEY OUTPUT

Agentic GBS charter on a page; ten dimensions answered, KPIs agreed, 18-month runway approved.

*Without it, every later decision gets re-argued.*

## 02 Discovery via *task mining*

A fast route to **real, evidenced AI opportunities** - without waiting on process documentation that doesn't exist or isn't current.

- **Skips the "document the process first" trap** that kills six months of any programme.
- **Sees the work as it actually runs** - including the 30% outside any SOP, in spreadsheets and email.
- **Surfaces AI-shaped patterns directly**: copy-paste loops, system-switching, repeat queries - each mapped to an agent design.
- **Quantifies the prize per opportunity** in effort and exception rate before any business case.
- **Use-case shortlist in 8-12 weeks**, not six months of workshops.

KEY OUTPUT

A ranked list of evidenced AI opportunities with effort, exception data and an agent design for each.

*Self-reported maps are wrong. Mining is ground truth.*

## 03 Top-down with the *use case library*

Cross-reference task-mining findings against the **Proservartner library of 88 use cases**.

- Solutions named: P2P, O2C, R2R, FP&A, Tax, Treasury, Controls, Risk - not generic patterns.
- Each entry has a problem statement, agent design, dependencies and a "quick win vs build" rating.
- Patterns benefiting multiple owners get prioritised (e.g. master data, reconciliations).
- Output: a prioritised pipeline of 15-25 with sequencing, approach and owner per item.

KEY OUTPUT

A prioritised pipeline of 15-25 use cases with owners and dependencies mapped.

*Bottom-up finds patterns; the library names solutions. Need both.*

## 04 Pilot *quick wins*

Pick **2-3 narrow, contained** use cases. Clean data, willing users. Time-box 8-12 weeks.

- Typical first wins: invoice reading, cash apps, ticket triage, contract extraction.
- Hybrid team: Proservartner + client + AI partner where appropriate.
- Instrument day 1: human-in-the-loop, accuracy tracking, override rates.
- Exit criteria: accuracy threshold met, override trending down, sponsor sign-off.

KEY OUTPUT

Live agents on production with measured accuracy and a credible reference inside the org.

*Not about ROI. About seeing an agent behave well.*

SCALE & EMBED

# Steps 5-7 - *from pilot to operating rhythm.*

## 05 Scale *end-to-end*

Turn the pilot into a connected agent layer for **the whole cycle**. One agent is automation; orchestration is GBS.

- Sequence 4-7 agents per cycle (e.g. AP: ingest, code, dedupe, pay, audit, insights, query).
- Move from human-in-the-loop to human-on-the-loop where accuracy holds.
- Wire in always-on agents: audit, insights, vendor/customer queries.
- Cycle owner becomes accountable for the whole agent flow, not just task quality.

**KEY OUTPUT**

A live cycle running on orchestrated agents with measurable cycle-time, exception and quality wins.

*Scale teaches you what was missing in step 1.*

## 06 Industrialise *the platform*

Stand up the platform layer **once**, reuse across cycles. Audit stops being a blocker.

- Model gateway, orchestration, observability, guardrails - built once.
- Small Agent Ops team (3-6 FTE) maintains, improves, deprecates.
- Governance embedded in existing risk/audit/change committees, not parallel ones.
- Reusable agent patterns so the 11th cycle starts from a library, not a blank page.

**KEY OUTPUT**

Agent #11 costs a fraction of agent #1, and audit can sign off.

*Without it, every new agent is a snowflake.*

## 07 Embed & *evolve*

Move from project to BAU. **The workforce shifts** from operators to reviewers.

- Task mining runs continuously; use-case pipeline refreshes quarterly.
- Capability uplift as a programme: prompts, evals, debugging, agent design.
- Adjacent cycles ship as planning decisions, not business cases.
- Operator-to-reviewer transition is structured, not improvised.

**KEY OUTPUT**

Next 10 agents ship without a steering committee.

*"Orchestrating agents" is a state, not a project.*

## All 7 The *quiet killers*

Four things that **thread through every step**. Under-invest and the programme stalls quietly, late.

- **Data** - fragmented master data makes every agent brittle.
- **Change** - operator-to-reviewer is the biggest shift in a decade.
- **Governance** - designed in at step 1, not bolted on at step 6.
- **Sponsorship** - programmes that lose their exec sponsor mid-flight don't recover.

**WHY THESE ARE QUIET KILLERS**

None kill on day one. They drain credibility over twelve months.

*The technology is the easy part.*

USE CASES

# Twelve agents that *earn their keep first.*

P2P
**Invoice** *ingestion*
Reads PDF/email invoices, codes to GL and cost centre, flags exceptions for review.

P2P
**Duplicate** *payment guard*
Cross-checks new invoices against history; catches near-dupes humans miss.

O2C
**Cash** *applications*
Matches remittances to open AR, handles short-pays and over-pays, escalates real disputes.

O2C
**Credit &** *risk review*
Real-time credit decisioning; dynamic limit reviews on portfolio shift; trade-finance signals.

R2R
**Reconciliation** *accelerator*
Bank rec, balance-sheet rec, GR/IR clearing - auto-matched with a confidence score per line.

R2R
**Close** *orchestrator*
Tracks close tasks, predicts slippage, escalates ahead of deadline, replaces the weekly close call.

FP&A
**Forecast** *first draft*
Reads actuals + drivers + commentary; produces the first-cut forecast for review and override.

CONTROLS
**SoD &** *controls testing*
Continuous testing of segregation, approval limits, journal patterns; flags exceptions to audit.

TREASURY
**Cash** *positioning*
Aggregates balances, predicts daily position, optimises sweeps and funding calls within policy.

TAX
**Statutory** *reporting*
Drafts statutory filings from the ledger; explains differences vs management view in narrative.

RISK
**Contract** *intelligence*
Reads customer and supplier contracts; surfaces volume-based discounts and revenue leakage.

FIN OPS
**Vendor** *query*
Answers vendor and customer email queries (status, copy invoice, dispute) with escalation.