

The Solutions24x7 AI Strategy (Starter Edition)

TCS FRAMEWORK · AI STRATEGY

STARTER EDITION

Document Overview: The pragmatic, fast-track AI Strategy for growing SMEs. It strips away enterprise complexity to focus purely on high-impact SaaS integration, AI literacy, and operational foundations.

1. EXECUTIVE SUMMARY

1.1 Strategic Context: Why AI Now

The AI adoption window is compressing. Organizations that establish governed AI operations now will build compounding advantages in cost efficiency, speed-to-market, and institutional knowledge that late movers will struggle to replicate. The risk has inverted: it is no longer the risk of adopting AI early — it is the risk of arriving late to a transformed operational baseline while competitors have already embedded AI into their core workflows.

2. BUSINESS CONTEXT AND STRATEGIC INTENT

2.1 Corporate Strategy & Operating Model Baseline

Before AI can be directed at business outcomes, the existing operating model must be mapped: core revenue model, primary workflows from intake to delivery, current technology stack, and performance baselines. This ensures every AI initiative traces to a defined business capability — not a disconnected tool purchase.

2.4 The "Chat+" Imperative: Beyond Standalone Tools

*Many organizations begin their AI journey by purchasing standalone licenses (e.g., Microsoft Copilot, ChatGPT Enterprise). While this solves the **Intake Layer**, it leaves critical gaps in governance, routing, and observability.*

If you are already using standalone chat tools:

"You have the Intake Layer sorted. But without an underlying Operating Model, you have an un-audited, un-governed tool being used casually by untrained staff, locked into a single vendor's pricing. Our strategy implements the routing, observability, automated security, and staff certification required to turn that casual tool into a measurable, secure, enterprise-grade operating system."

3. AI VISION, PRINCIPLES, AND GUARDRAILS

3.3 Explicit Guardrails: Where AI Will NOT Be Used

Guardrails are codified directly into the Rosetta Strategy Config — not advisory statements. Examples: AI will not be used for direct customer communication without a Human-in-the-Loop review step. No agent will have write-access to financial or HR systems without an accountable human owner. PII will not pass to external LLM endpoints without tokenisation at the gateway layer. AI outputs used in compliance contexts must carry a disclosure label.

4. CURRENT STATE AND AI READINESS

4.2 Process and Workflow Baseline

Workflows are mapped across three categories: **Fully automatable** (structured, high-volume, deterministic tasks); **AI-augmented** (knowledge-intensive tasks where AI accelerates but a human validates the result); and **Human-only** (accountability-critical or outside the declared risk appetite). This mapping produces the "Functional Spine" — the prioritised sequence of workflows the Rosetta OS will govern in Phase 1.

4.3 Existing "Shadow AI" and Unsanctioned Tool Usage

Shadow AI — unsanctioned use of tools like ChatGPT, Claude, or Gemini by employees — is already present in most organisations before the formal strategy is written. Rather than suppressing it (which drives it underground), the assessment documents it: which tools, which teams, what data is being shared. This inventory becomes the business case for a governed Intake layer: a sanctioned, secure alternative that removes the incentive to circumvent governance while preserving the productivity benefit.

5. OPPORTUNITY LANDSCAPE & STRATEGIC INITIATIVES

5.1 Value Framework (Revenue, Cost, Experience)

Every AI investment maps to one of three value levers: **Revenue** (new capability, faster delivery, AI-powered client services); **Cost** (hours saved, error reduction, vendor consolidation); or **Experience** (client satisfaction, staff burden reduction, response speed). For each initiative, a value hypothesis is registered at inception and measured at 30, 90, and 180 days post-deployment — preventing AI spend from becoming undifferentiated "innovation budget".

5.3 Prioritized Portfolio: Quick Wins

Phase 1 initiatives are selected using a value-vs-complexity matrix: **Quick Wins** (high value, low complexity) activate first to demonstrate ROI within 90 days and build organisational confidence. Every initiative requires a documented value hypothesis, an assigned HITL owner,

and a measurable 90-day success criterion before entering the portfolio.

6. TARGET ARCHITECTURE AND TECHNOLOGY STACK

6.3 AI Ecosystem and Vendor Landscape Alignment

Rosetta Alignment: *The architecture leverages the foundation of the "Killer AI Stack":*

• **Intake:** Secure User Interface • **Strategy/Routing:** LLM Gateway & Routing Layer

8. PEOPLE, HYBRID WORKFORCE, AND CULTURE

8.3 AI Workforce Readiness (The "Rosetta Ready" Program)

The Rosetta Ready program certifies staff across three tracks: **Operator** (triage and dispatch — the Concierge and Pilot roles); **Governor** (review, authorise, and audit — the Captain roles); and **Builder** (configure and maintain agents). Certification confirms staff can safely occupy the 5 Human-in-the-Loop roles before any agent activates. The program runs in parallel with Phase 1 infrastructure deployment — not as an afterthought.

Workforce readiness is assessed before Phase 1 completion, not after. The Rosetta Ready diagnostic maps each team member to the appropriate HITL role, identifies digital-confidence gaps, and prescribes a targeted certification track. This prevents the most common AI governance failure: deploying capable agents to teams without the training or authority to supervise them.

9. IMPLEMENTATION APPROACH AND ROADMAP

9.1 Phased Execution Principles

The governing principle is **prove before you scale**. No capability advances to the next phase without meeting measurable success criteria at the previous stage. The Rosetta maturity journey: **Lite** (secure intake + first governed agent); **Ready** (OS active across primary workflows, staff certified, observability live); **Core** (automated orchestration, active governance dashboards); **Full** (embedded operating model with continuous improvement loops and standard agent commissioning).

9.2 Phase 1: Foundations & Quick Wins (0–6 Months)

- **Core Infrastructure:** Deploying a secure internal chat interface and an LLM routing gateway.
- **The Concierge Agent:** Launching a primary routing agent to act as the digital "front door" for staff.
- **Prompt Integration:** Embedding proven prompt frameworks seamlessly behind the interface.
- **Agent Register:** A governed directory of internal-only AI assistants.

9.3 Phase 2: Scale & Governance (7–18 Months)

Phase 2 extends the Rosetta OS to secondary workflows and deploys specialist AI agents for core business functions. The SteerCo governance cadence is established — quarterly reviews against KPI baselines, active observability dashboards replacing manual monitoring, and a formal agent commissioning process replacing ad-hoc deployment. Phase 2 budget is gated by Phase 1 ROI evidence, documented at the 180-day review.

11. MEASUREMENT AND CONTINUOUS IMPROVEMENT

11.1 KPI Framework (Business + Technical Metrics)

Business KPIs: Time-to-complete for targeted workflows vs pre-AI baseline; error rate on AI-assisted outputs; staff hours reallocated from automated tasks to higher-value work; AI adoption rate by department.

Technical KPIs: Triage Accuracy (correct routing on first pass); Coordination Velocity (intake-to-output resolution time); System Uptime (target 99.5%+); Trace Coverage (% of LLM interactions captured); Gateway Error Rate (policy violations per 1,000 requests).

💡 What's in the Full Version?

The Definitive Master Strategy includes 12 comprehensive sections — deep enterprise components skipped in this Starter Edition:

- **Data & Infrastructure Readiness (Sec 4.1):** Assessing silos, cloud setup, and pipeline maturity.
- **AI Operating Model & Governance (Sec 7):** Full "Rosetta" framework, Squads, and automated PII Risk Management.
- **AgentOps & Hybrid Work Design (Sec 8.1, 8.2):** Managing autonomous digital workers and explicit human/AI handoffs.
- **Financials, Investment, & ESG (Sec 10):** Budgeting models, granular ROI tracking, and environmental compute metrics.
- **Continuous Eval & SteerCo (Sec 11.2, 11.3):** Active Observability dashboards and structured executive review cadences.
- **Appendices (Sec 12):** Vendor evaluations, acceptable use policies, and standing orders.

The 15-page Definitive AI Strategy is included in our comprehensive AI Framework Engagement — solutions24x7.com