



TAYCAN^{AI}



The AI Deployment Playbook

How Taycan AI Takes You From Problem to Working System.

INTRODUCTION

Most AI projects fail. Not because the technology does not work. Because the project was scoped wrong.

Industry research consistently finds that the majority of enterprise AI projects never make it into production. Of the ones that do, a significant share are quietly abandoned within the first year. The root causes are consistent: scope that was too broad, data that was not accessible, systems that were never properly integrated, and teams that were not brought into the process.

This is not a technology problem. It is a deployment problem.

Taycan AI was built to solve it. We deploy focused, integrated AI systems that handle specific operational workflows. We do not sell roadmaps or strategy decks. We build working systems, measure what they deliver, and expand from there.

This playbook describes exactly how we do it. Four phases, in order, with clear deliverables at each stage and defined success criteria before anything goes live.

What This Playbook Covers

Phase 1 – Discover: Map the workflow, identify the data, define the baseline metric.

Phase 2 – Architect: Design the AI system and plan every integration before building.

Phase 3 – Deploy: Build, validate against real data, and go live with full monitoring.

Phase 4 – Measure: Track results against baseline, refine, document the ROI.



Our Commitment: Every Taycan AI engagement ends with a working system in production, not a presentation. We scope each project tightly, deploy quickly, and measure the result against a clear baseline established before the first line of code is written.

PHASE 1: DISCOVER

Why discovery determines everything.

The most common failure mode in AI deployment is starting to build before the workflow is properly understood. Teams underestimate the complexity of the data, miss key system dependencies, or discover halfway through build that the information the agent needs is locked in a system that cannot be accessed.

Discovery eliminates this. We spend one to two weeks mapping the workflow in full before any architecture decisions are made. The output of this phase is a clear, validated picture of what we are building, what data it will use, and what success will look like.

What Happens in Discovery

Workflow mapping: We document every step of the target workflow — inputs, outputs, decision points, exceptions, and edge cases. If someone currently does this work manually, we interview them. We map what they actually do, not what the process document says they do.

Data audit: We identify every data source the workflow touches. We verify accessibility — does it have an API, an export function, or a database connection we can reach? We flag any data quality issues that need to be resolved before the agent can work with them.

System inventory: We list every system involved. CRM, ERP, helpdesk, data warehouse, spreadsheets. We document the integration method for each and confirm that the connections are technically feasible within your environment.

Baseline measurement: We establish the current state in measurable terms before anything changes. How many hours does this workflow take per week? What is the

current error rate? What does a good output look like? This baseline becomes the measure of ROI after deployment.

Scope confirmation: We confirm the scope of the first deployment. Not everything will be automated in the first phase. We identify the highest-value, most technically feasible starting point and define clear boundaries for what the first system will and will not handle.

Phase 1 Deliverable: Workflow map, data audit report, system inventory, baseline metrics, and confirmed scope. Typical duration: 1 to 2 weeks depending on workflow complexity and number of systems involved.

PHASE 2: ARCHITECT

Design first. Build second. Always.

Architecture is where most AI projects skip ahead. The tooling is available, the team is enthusiastic, and it is tempting to start building. The result is a system that works in isolation and breaks when it meets the real operational environment.

We do not start building until the architecture is complete and agreed. Every integration is planned, every human touchpoint is defined, and every edge case has a documented handling approach. The architecture document is a contract, it defines exactly what will be built and how it will behave before construction begins.

What the Architecture Covers

Agent logic design: We define the decision rules the AI will follow at each step. Classification criteria, confidence thresholds, output formats, and the exact behaviour for every defined input type. Nothing is left to interpretation at build time.

Integration architecture: We design the connection layer between the AI system and your existing tools. Every API endpoint, authentication method, data transformation, and error handling path is documented before build begins. We connect to your systems – we do not replace them.

Human touchpoint design: We define exactly where a human is in the loop. Which outputs require review before action? Which can be executed automatically? What does the review interface look like? Human oversight is a design decision, not an afterthought.

Edge case handling: We document every known edge case and define the system behaviour for each. Low-confidence results, missing data, system outages, and unusual inputs all have defined handling paths. The agent knows what to do when things do not go to plan.

Validation framework: We define how the system will be tested before going live. Which historical data will be used? What accuracy threshold must be met? Who signs off on the validation results? These criteria are agreed before build, not after.

Phase 2 Deliverable: Technical architecture document, integration plan, edge case register, and validation criteria. Typical duration: 1 week. Architecture is agreed in writing before Phase 3 begins.

PHASE 3: DEPLOY

Nothing goes live until it has been validated.

Build phase is where the system is constructed, integrated with your environment, and tested against your actual operational data. We do not validate against synthetic data or invented scenarios. We run the system against real historical cases and compare its output to how your team actually handled them.

Before full deployment, the system runs in parallel with your existing process. Your team sees the AI output alongside the normal workflow. This builds confidence, surfaces edge cases in the real environment, and allows calibration before the system takes over.

The Four Steps of Deployment

Step 1 – Build The system is constructed according to the architecture document. Integrations are connected and tested against staging environments. Edge case

handling is implemented and verified. The build is complete when every component behaves as specified in the architecture.

- All integrations connected and tested in staging environment
- Edge case handling implemented and verified against the register
- Internal QA completed before external validation begins

Step 2 – Historical Validation The system is run against 60 to 90 days of real historical data. Every output is compared to the actual outcome – how did your team handle that case? Where does the AI agree? Where does it differ? Accuracy is benchmarked against the threshold defined in Phase 2.

- 60–90 days of real operational data used for validation
- Output accuracy benchmarked against agreed threshold
- Discrepancies reviewed and system calibrated accordingly

Step 3 – Parallel Run The system runs alongside your existing process for one to two weeks in the live environment. Your team reviews AI output in real time without the system taking action yet. This surfaces the edge cases that historical data did not reveal and builds team confidence.

- System runs live but does not take action without review
- Team feedback collected daily and incorporated
- Parallel run continues until accuracy threshold is consistently met

Step 4 – Go Live The system takes over the workflow according to the defined human touchpoint design. Monitoring is active from the first day. Any output that triggers an escalation threshold is flagged automatically. The operations owner receives a weekly summary for the first 30 days.

- Automated monitoring active from day one
- Escalation thresholds defined and tested before cutover
- Weekly summary report to operations owner for first 30 days

Phase 3 Deliverable: Working system in production, monitoring dashboard, and go-live sign-off report. Typical duration: 2 to 6 weeks depending on integration complexity and number of systems involved.

PHASE 4: MEASURE

The baseline established in Phase 1 becomes the measure of value.

An AI system that cannot demonstrate its value in measurable terms is a liability, not an asset. We establish the baseline before deployment so there is no ambiguity about what changed and why. Every metric tracked in Phase 4 is compared directly to the state we documented in Phase 1.

AI systems also improve with feedback. The Measure phase is where that feedback is collected, reviewed, and acted on. Edge cases that emerge in production are handled. Accuracy drifts are identified and corrected. The system gets better over time because there is a structured process for making it better.

What We Track

Output volume: Total workflow handled by AI vs. manually per week, compared to pre-deployment baseline. Time saved: Labour hours returned to the team per week. Calculated against the baseline established in Phase 1. Accuracy rate: Percentage of AI outputs that required no correction or override. Tracked weekly, reported monthly. Error rate: Edge cases and failure modes logged, categorised, and resolved. Trend tracked month over month. Escalation rate: Percentage of cases flagged for human review. Should trend down as the system is refined. ROI: Labour cost displaced versus total system cost. Calculated quarterly with full methodology disclosed.

The Monthly Review

Every month we review the system performance with the operations owner. We look at accuracy trends, edge case patterns, escalation rates, and any changes in the underlying workflow. If refinements are needed, they are scoped, prioritised, and scheduled. The monthly review is also where expansion opportunities are identified – once one workflow is running well, the next highest-value opportunity becomes the natural next conversation.

Phase 4 Deliverable: Monthly performance report, ROI calculation, and refinement log. Ongoing.

Duration: Ongoing. Measure phase runs for the life of the deployment.

NEXT STEPS

Ready to deploy AI in your operation?

The first step is a 30-minute strategy call. No pitch. No deck. A direct conversation about your workflows and what AI can do for them.

What Happens on the Strategy Call

You describe your operation. We listen for workflow patterns, system dependencies, and data availability. We ask specific questions about volume, frequency, and current process.

We identify the opportunity. Based on what you tell us, we identify the highest-value workflow to automate first. We explain why and what the deployment would involve.

We outline the deployment. We walk through what Phase 1 through 4 would look like for your specific situation. Realistic timeline, realistic scope, realistic outcome.

You decide what to do next. No obligation.

Book your free 30-minute strategy call at www.taycanai.com

“

The future of business is AI-first.

Satya Nadella, CEO, Microsoft

”