

Playbook

From hype to impact

Turning AI into enterprise value



How to use this playbook

This playbook is designed for executives, transformation leaders, CIOs, CFOs, HR directors, procurement, operations, ITSM, and shared services leaders who recognise AI's potential but need to decide where to start.

Use it when your organisation has moved beyond curiosity and requires a clear path to action.

How to use it:

Work through in order, or jump to the section that matches the stage you are at:



Diagnose where value is leaking from existing business systems →



Build an initial use case backlog →



Score and prioritise your first use case →



Establish governance for production →



Optimise AI capabilities with Agentic Core →

A proven playbook

As demonstrated by the UWE support agent (detailed in our eBook), which now handles 80% of routine queries, these approaches help identify your own high-impact starting point.



Diagnose where value is leaking

The strongest opportunities sit between core systems.

ERP, ITSM, CRM, HR, finance, and procurement platforms work well for standard processes, but people still spend significant time chasing information, correcting errors, resolving exceptions, and coordinating across teams.

Value leak	What it looks like	Where agents can help
Process friction	Status chasing, manual clarification, approvals stuck in email, and users unclear on next steps.	Gather context, recommend actions, draft updates, route work, and maintain case notes.
Exception overload	A small proportion of cases consumes disproportionate effort.	Classify exceptions, check policy, retrieve evidence, resolve safe cases, and escalate what needs judgement.
Rework and duplication	Same data entered, checked, copied, and corrected multiple times.	Validate info, compare records, flag mismatches, and update permitted systems.
Delay costs	Suppliers, employees, students, customers, or service users wait for manual handoffs.	Reduce cycle time with easier, faster, consistent next actions.
Control gaps	Rules exist, but they're applied inconsistently or audited after the fact.	Apply policy at the point of work and leave a trace of decisions, evidence, and escalation.
Experience damage	People feel they are waiting for the organisation rather than being served by it.	Provide instant guidance, self-service support, and smooth handover to human teams.



Mark the two or three rows above that describe your organisation most accurately.

These are your hunting grounds.



Build the first use case backlog

Target workflows with visible pain points, recurring volume, clear ownership, accessible data, and measurable KPIs.

A strong first use case demonstrates that agentic AI can improve real work without destabilising core systems.

Core use-cases		
Area	Potential use cases	Metrics
Finance	Invoice exceptions, reconciliation, close preparation, payment queries, and policy checks.	Minutes saved per case, exception backlog, error rate, close-cycle time, and working-capital impact.
ITSM	Ticket triage, knowledge retrieval, password/access support, outage updates, and escalation routing.	Query deflection, MTTR, backlog, first-contact resolution, and employee experience.
HR	Employee support queries, onboarding, policy questions, case handling, and knowledge access.	Query response time, HR workload, accuracy, escalation rate, and employee satisfaction.

Extended use-cases		
Area	Potential use cases	Metrics
Procurement	Vendor onboarding, approval routing, supplier follow-up, policy checking, and contract queries.	Cycle time, supplier activation speed, approval delay, compliance, and user satisfaction.
Operations	Order exceptions, service updates, workflow coordination, and cross-system reconciliation.	Throughput, delay reduction, rework, service quality, and customer experience.
Customer/service	Self-service agents, assisted support, guided journeys, and document handling.	Containment, conversion, retention, satisfaction, and cost-to-serve.



Pick five to ten candidate use cases across two or three functions.
Don't try to be exhaustive – the scorecard should do the filtering.

Score and prioritise the first use case

A great first use case is usually where pain is visible, where the agent can be bound safely, and where outcomes are measurable.

Use this scorecard to compare options:

Criterion	Question to ask	Score (1-5)
Business pain	Is the problem visible to leaders/users?	
Volume	Does it happen often enough for material gains?	
Boundary fit	Repeatable/bounded but not fully automated?	
Manual coordination	Slowed by chasing/interpretation/routing/follow-up?	
KPI clarity	Measurable in time/cost/quality/risk/experience/revenue?	
Reusability	Creates patterns/connectors/governance for next cases?	
Data readiness	Can the agent access the needed knowledge/systems/context? <i>Note: This is a dealbreaker category; if you score a 1 here, you should disqualify the use case for now and revisit it.</i>	
Risk manageability	Safely bound with human oversight? <i>Note: This is a dealbreaker category; if you score a 1 here, you should disqualify the use case for now and revisit it.</i>	

Scoring Guide:

32-40: Strong first use case. Build the business case.

24-31: Promising. Tighten data access or governance before committing.

Below 24: Park it. Re-score in 6 months or pick another candidate.

A high-scoring use case still requires clear ownership, defined data access, a governance framework, a value case, and a realistic path to production.

However, the scorecard helps prevent your agentic programme from being driven by novelty.



Govern for production

The more an agent can do, the clearer the control model needs to be.

This is especially important when agents operate within enterprise environments and connect to systems containing sensitive operational, employee, customer, or financial data.

Each row below needs a named decision-maker before the agent goes live:

Governance area	Decision required	Decision Maker
Data access	What data does the agent need, and how is access limited by role, task, and purpose?	
Action rights	Can the agent read, recommend, draft, update, approve, or transact? Which actions need human approval?	
Human oversight	Which cases escalate because they are high risk, low confidence, regulated, customer-sensitive, or financially material?	
Auditability	What trace is left showing what the agent saw, what it did, why it did it, and who approved it?	
Evaluation	How will accuracy, policy adherence, exception handling, experience, and business impact be measured over time?	
Ownership	Which business process owner is accountable for the outcome, not just the technology?	

Back to top →

Agentic Core - a repeatable AI delivery model

The first agent proves value, the second is cheaper, the third is faster.

Agentic Core is the delivery model that makes compounding possible.

Carrying AI adoption into production, then beyond go-live into managed continuous improvements.

Each stage produces an output that feeds the next, and each deployment leaves behind reusable governance, patterns, and integrated learning.


Stage	What happens	Output
1. Ideate	Identify pain points, opportunities, users, and business outcomes.	Long list of candidate use-cases.
2. Prioritise	Score use-cases by value, feasibility, risk, data readiness, and reusability.	Prioritised backlog and first use case recommendation.
3. Shape value	Define KPIs, benefit logic, cost assumptions, ownership, and decision gates.	Business case and delivery brief.
4. Design	Map workflow, data, tools, action rights, escalation, evaluation, and governance.	Solution design and control model.
5. Build	Develop the agent, integrations, prompts, tools, tests, and user experience.	Working agent in a bounded environment.
6. Deploy	Release in a controlled way with monitoring, user feedback, and adoption support.	Production deployment and initial value evidence.
7. Manage and optimise	Monitor performance, evaluate outcomes, improve the agent, and manage change.	Ongoing service improvement and next use case readiness.



AgenticCore brings together structured ideation, business-case discipline, solution design, controlled deployment, managed oversight, and continuous optimisation - the delivery model behind every agent we help organisations put into production.



Take the first step



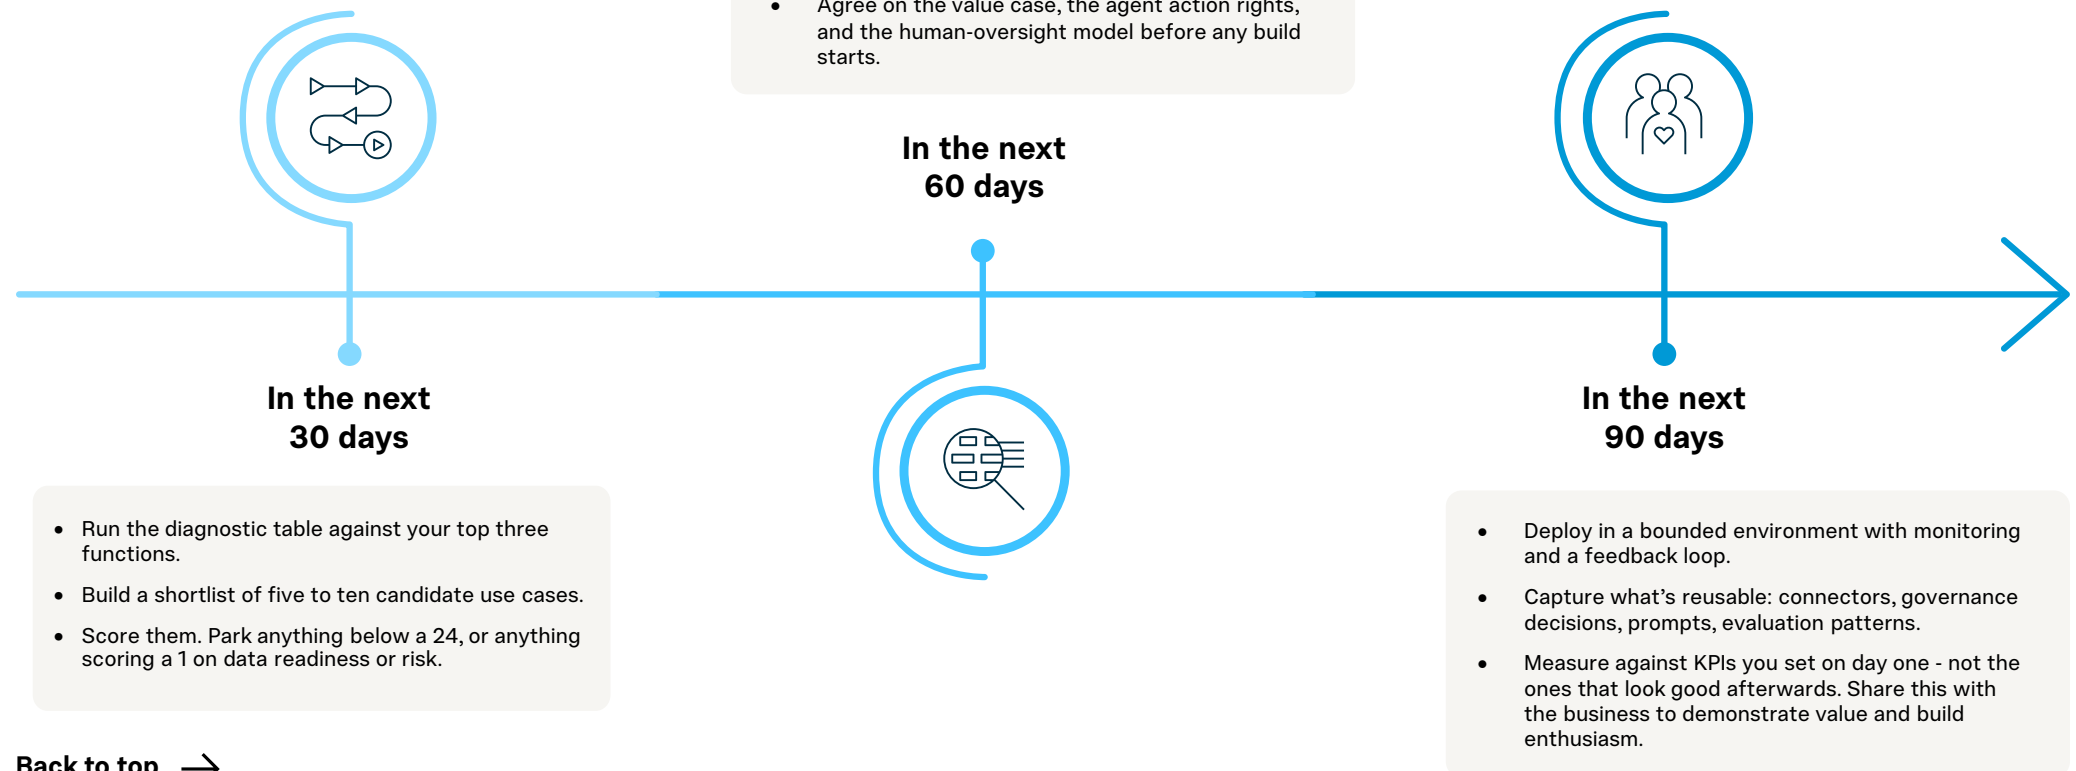
The last point is the one most leaders pause on.

It also reveals some compelling areas that may not have been considered.

Back to top →

Your next 90 days

This playbook works best when it ends with a calendar entry.



Back to top →

If you'd like help running
the playbook...

...we'd be glad to scope your first use
case with you.

Start the conversation →



salesenquiries@nasstar.com

0345 003 0000



Navigate complexity, with confidence

nasstar.com