Kyndryl - Agentic AI Automation for ITSM & Security Operations

Proposal by ThirdEye Data AI

Prepared for: Kyndryl

Table of Contents

1 Executive Summary

AI Ops is an agentic AI platform for IT service management and security operations. It reads every incoming ticket, determines the correct course of action by reference to how Kyndryl's own L1 and L2 engineers have resolved comparable cases, executes the approved steps directly against Kyndryl's connected systems, and writes the outcome back into Kyndryl's system of record — with full audit logging and a human checkpoint on anything consequential.

This is not a chatbot placed in front of a service desk, and it is not a static runbook engine. It is a decision-and-execution platform, configured to Kyndryl's environment, it's task categories, towers, and risk policies.

What is being proposed

- **An AI Ops platform,** configured to Kyndryl's ticket taxonomy, its assignment structure and approval policy — not a generic product installed unchanged.
- **A model trained on Kyndryl's own resolution history,** so that every decision the platform makes is grounded in evidence drawn from Kyndryl's environment rather than generic assumptions.
- **An execution layer of 24 connectors** across identity, network and security, cloud, VDI and endpoint, collaboration and SIEM.
- **Deployment in Kyndryl's own infrastructure,** so that code, models, ticket data and the audit trail remain inside its environment throughout, with no vendor-hosted component in the path of a ticket.
- **A governance layer** in which autonomy is granted category by category, on evidence, and never assumed.
- **A development and customisation programme** of 100+ features delivered across three milestones of four weeks each *(subject to the findings of the Due-Diligence phase)*, sequenced by build complexity and each closing with a production rollout and user acceptance testing.

The delivery shape

Milestone	Weeks	Complexity	Theme	Features
Due Diligence & Execution Planning	Week 1-2		1 – Detailed due diligence into Kyndryl's environment, business processes, data landscape, existing tickets, etc. 2 – Develop a detailed project plan for implementation of the prioritized features list along with execution timelines, resource allocation, and cloud infrastructure needed.	
Milestone 1	Weeks 3 - 6	Easy	Foundation and Core Automation	44
Milestone 2	Weeks 7 - 10	Medium	Assisted Execution and Governance	39
Milestone 3	Weeks 11 - 14	Complex	Advanced Operations and Autonomy	17
Total	14 weeks	—	Full platform in production	100

Features are ordered by build complexity — the straightforward, high-volume capability is delivered first, the medium-complexity governance and execution depth second, and the advanced capability last. Each milestone ends with a rollout into production and a UAT sign-off gate, so value is in use from week four rather than at the end of the program.

A residual share of work remains human-owned by design. Physical hardware dispatch, live security incident response, privileged access grants and irreversible data operations are excluded from automation as a matter of policy, not capability.

How the engagement is structured

ThirdEye Data delivers the design, development, customisation, integration and deployment of the platform. Kyndryl will provide the infrastructure (cloud or bare-metal) it runs on, in its own accounts, and meet the infrastructure cost directly. Development and customisation are charged per milestone against acceptance; there is no licence fee, subscription fee or per-ticket charge. Section 15 sets out estimated infrastructure and GPT API cost, and Section 16 the commercial structure.

2 Problem Statement

IT operations at enterprise scale carry a cost structure that rises in direct proportion to ticket volume. Every additional ticket consumes engineer time; every increment of volume growth therefore requires an increment of headcount. The constraint is structural, and no amount of process optimisation removes it.

2.1 Operating context

Dimension	Position
Annual ticket volume	Approximately 200,000 tickets across incident, request and security queues
Monthly run rate	Approximately 16,700 tickets per month
Security and access load	Approximately 4,000 tickets per month
Escalation rate	Approximately one ticket in three escalates from L1 to L2
Coverage requirement	Continuous 24x7 cover across multiple towers
Task taxonomy	Eighteen distinct task groups spanning access, endpoint, software, security and network

2.2 Six structural problems

Problem 1 — Cost scales linearly with volume

Because resolution effort is manual, cost per ticket is effectively fixed. Doubling volume doubles cost. There is no mechanism by which the estate becomes cheaper to run as it grows, and no point at which prior investment reduces marginal cost.

Problem 2 — Escalation multiplies effort

Roughly a third of tickets escalate from first line to second line. Each escalation adds queue wait, context hand-off and a degree of rework, and consumes the time of the most expensive engineers on the team. The escalated population is also the population most likely to breach service targets.

Problem 3 — Resolution knowledge is undocumented and concentrated

The most reliable knowledge about how a given failure is actually resolved exists in the working memory of experienced engineers and in the free-text closure notes of past tickets. It is not systematically captured, it does not transfer on attrition, and it is not available at the moment a new ticket arrives.

Problem 4 — Risk accumulates while security tickets queue

Security and access work carries consequence in the waiting. A delayed access revocation, an unapplied geo-blocking rule or a deferred firewall change during an active incident each widen exposure. Queue time on this class of ticket is not merely a service inconvenience; it is a risk position.

Problem 5 — Continuous coverage is structurally expensive

Sustaining a single seat on a 24x7 rota requires approximately five engineers once leave, training, attrition and shrinkage are accounted for — materially more than a flat working-week calculation suggests. Coverage cost is therefore incurred against the clock rather than against demand, and idle capacity at low-volume hours cannot be recovered.

Problem 6 — Nothing compounds

Each ticket is solved substantially from first principles. The thousandth password reset costs what the first one cost. Effort expended on resolution produces no durable asset, so the operation never accrues leverage from its own history.

2.3 Why additional headcount does not resolve this

Adding engineers addresses capacity but preserves the cost curve. It also compounds the knowledge concentration problem: a larger team increases variance in how comparable tickets are handled, which in turn increases rework and weakens the consistency of the audit position. The requirement is not more capacity applied to the same model. It is a different model.

The objective is to convert resolution history into an operating asset — so that work already done reduces the cost of work still to come.

3 Solution Approach

3.1 What is ThirdEye Data providing

The engagement covers the design, development, customisation, integration and deployment of an AI-driven operations platform, built on a proven baseline and running on Kyndryl's own infrastructure. It is delivered as a working system rather than assembled from scratch, and the development and customisation effort is directed at four things:

1. **Kyndryl's taxonomy** — Kyndryl's task groups, categories and subcategories are loaded so that classification reflects Kyndryl's taxonomy rather than a vendor's.
2. Kyndryl's **structure** — Kyndryl's towers, assignment groups and queue owners are mapped so that every ticket is routed to the team that actually owns it.
3. Kyndryl's **policy** — risk tiers, approval requirements, deny-lists and blast-radius limits are set to Kyndryl's policy, and is upto Kyndryl to change.
4. Kyndryl's **systems** — connectors are pointed at Kyndryl's instances of the systems Kyndryl already runs, using each vendor's native API and CLI.

The platform is then deployed into Kyndryl's environment and operates continuously against Kyndryl's live ticket flow. Section 14 sets out the infrastructure it runs on, and Section 16 how the engagement is structured commercially.

3.2 AI is the backbone, not a feature

Artificial intelligence in this platform is not an assistive layer bolted to a workflow engine. It is the mechanism by which the system decides what to do. Four capabilities carry the platform, and each is derived from Kyndryl's data:

Capability	What it does	Why it matters
Classification	Determines category, subcategory, priority and owning tower for every incoming ticket	Removes the manual triage step and makes correct routing automatic from the first second
Precedent retrieval	Finds the nearest comparable historical cases and the applicable knowledge article	Grounds every decision in what Kyndryl's engineers actually did, with citable evidence
Planning	Produces an ordered, executable action plan across one or more connected systems	Converts a diagnosis into a specific sequence of commands rather than a recommendation
Learning	Converts every human approval, edit, rejection and override into a training example	The platform improves from use, so the automated share rises over time without redevelopment

The consequence is that the platform gets better at Kyndryl's estate specifically. A generic model applied to Kyndryl's tickets would encode assumptions from other organisations. A model built on Kyndryl's closure history encodes Kyndryl's standards, it's naming conventions, approval customs and architecture.

3.3 What differentiates this approach

It executes — it does not only advise

The majority of tools in this category classify, summarise or suggest, and leave a human to perform the work. This platform completes the action against the target system and closes the ticket. The unit of value is a resolved ticket, not a recommendation.

Every decision carries its evidence

Each proposal is presented with the knowledge article it relies on, the historical tickets it matches, the success rate of that resolution path and a confidence score. A reviewer is never asked to trust an unexplained output.

Autonomy is earned, category by category

No category begins in autonomous operation. Each runs first in shadow, then assisted, and is promoted to zero-touch only once measured precision justifies it. Promotion is a governed decision supported by evidence, and it is reversible.

Guardrails are Kyndryl's **policy, held as configuration**

Risk tiers, mandatory-approval categories, deny-lists, blast-radius ceilings and change-freeze windows are configuration owned by Kyndryl's teams. They are not compiled into the product and do not require a vendor release to change.

Kyndryl's **system of record remains authoritative**

The platform is a system of action operating alongside Kyndryl's system of record. Ticket state, notes and outcomes are written back continuously, and the service management platform remains the single source of truth for reporting, audit and contractual measurement.

It is built on infrastructure Kyndryl **owns**

The platform is fully containerised and is deployed into Kyndryl's own data centre, dedicated hardware, or cloud subscription. Capacity is provisioned in accounts Kyndryl control, infrastructure cost is incurred directly by Kyndryl at Kyndryl's own negotiated rates, and the deployed platform continues to operate independently of any ongoing commercial relationship.

Kyndryl's data stays under Kyndryl's control

Classification and precedent retrieval run on models trained on Kyndryl's own history and hosted inside Kyndryl's environment; that data never leaves it, and neither does the ticket store or the audit log. The reasoning step calls a hosted GPT API, to which only the minimum context required for a decision is sent, with identifiers and credentials redacted before transmission and no retention for model training. Where policy requires it, that step can be moved to a self-hosted model.

The differentiator is not that the platform uses AI. It is that the AI is trained on Kyndryl's resolution history, constrained by Kyndryl's policy, and connected to Kyndryl's systems — which is what converts a model into resolved tickets.

4 Platform Architecture

The platform is organised into seven layers. Each layer has a defined responsibility and a defined interface to the layers above and below it, which allows components to be replaced, scaled or re-pointed without disturbing the rest of the system.

Figure 1 — Layered platform architecture

4.1 Layer responsibilities

Layer	Responsibility	Principal components
Presentation	Everything a human sees or acts through — monitoring, investigation, approval and reporting	Overview dashboard, Live Ops Console, Ticket Explorer, tower-lead views, approval workspace, reporting
Orchestration	Controls the lifecycle of every ticket and coordinates the specialist agents	Agent orchestrator, workflow engine, ITSM state machine, SLA engine, escalation and routing
Intelligence	Produces the decision — what this ticket is, what was done before, what to do now	Classifier, precedent retrieval over a vector index, planner, confidence scoring, learning pipeline
Governance	Constrains the decision before anything is permitted to execute	Guardrail policy engine, risk tiering, blast-radius calculator, rollback manager, approval broker
Integration	Carries the approved action to the target system using that vendor's native interface	24 connectors across identity, network and security, cloud, VDI and endpoint, collaboration, SIEM
Data	Holds the ticket record, the evidence base and the immutable trail	Ticket store, action-trace store, knowledge base, vector index, audit log, metrics store
Security & platform	Enforces identity, protects credentials and provides the runtime	RBAC mirrored from AD and Entra ID, secrets vault, encryption, container runtime, HA and scaling

4.2 Design principles

- **Separation of decision and execution —** no component holds credentials, evaluates policy and executes actions. The layer that decides is not the layer that acts.
- **Reversibility by default —** every action carries a rollback path, generated and validated before the action is offered for approval.
- **Least privilege —** the platform never holds broader privilege than the human role whose work it is performing.
- **Auditability at source —** the audit record is written as the action occurs, not reconstructed afterwards.
- **Fail safe, not fail open —** connector failure, model unavailability or policy ambiguity results in a hand-off to a human, never in an unlogged or partial action.

5 Data Flow

Every ticket follows the same five-stage path, whether it completes without human involvement or is handed to an engineer at the first checkpoint. The path is uniform; only the decision at stage three differs.

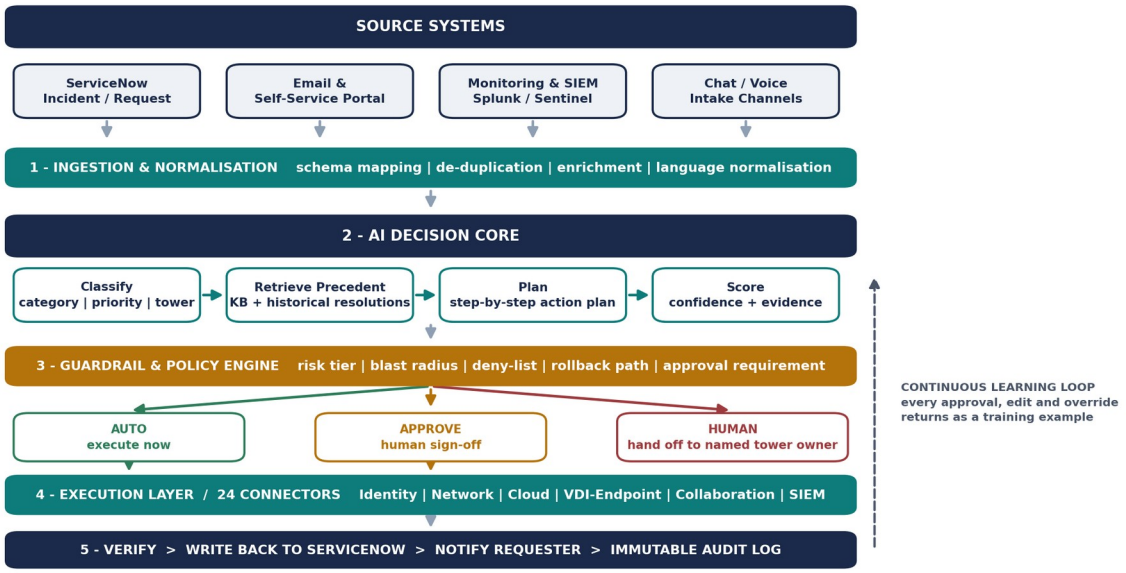

Figure 2 — End-to-end data flow

5.1 Stage by stage

Stage 1 — Ingestion and normalisation

Tickets are drawn from the service management platform in near real time, and can additionally be raised from email, a self-service portal, monitoring and SIEM alerts, or chat and voice channels. Incoming records are mapped to a common schema, de-duplicated against open work, enriched with requester, asset and entitlement context, and normalised for language and formatting.

- **Example —** a monitoring alert reporting disk utilisation above threshold on a production host is ingested, matched against an existing open ticket for the same host, and suppressed as a duplicate rather than raising a second ticket.

Stage 2 — AI decision core

The decision core executes four operations in sequence. It classifies the ticket to category, subcategory, priority and owning tower. It retrieves the nearest historical resolutions and the applicable knowledge article. It produces an ordered action plan across whichever systems are required. It scores its own confidence and attaches the supporting evidence.

- **Example —** a request for access to a finance reporting application is classified to Access and Identity / Application Access, matched to eleven comparable historical grants and one knowledge article, and planned as a two-step group addition with entitlement verification, at 96% confidence.

Stage 3 — Guardrail and policy engine

Before anything is permitted to execute, the plan is evaluated against policy. The engine assigns a risk tier, calculates blast radius, checks every step against deny-lists and change-freeze windows, confirms a

tested rollback path exists, and determines whether approval is required. The output is one of three routes: execute automatically, execute after human sign-off, or hand off to a named owner.

- **Example —** a proposed firewall rule change is found to affect 4,200 users and to fall inside a mandatory-approval category, so it is routed to the approve path with a full impact summary rather than executed.

Stage 4 — Execution

The approved plan is carried out step by step through the relevant connectors, each using the native API or CLI of the target system. Every command issued and every response received is logged. Steps execute in dependency order, failures halt the sequence, and the rollback path is invoked automatically where a partial execution would otherwise leave the estate inconsistent.

- **Example —** a VDI provisioning plan executes an Entra ID group addition, a Citrix delivery-group assignment and an Intune policy application in order, verifying each before proceeding to the next.

Stage 5 — Verify, write back and close

The platform confirms the intended end state has been reached, writes the outcome back to the service management platform as ticket state and resolution notes, notifies the requester, and seals the audit record. Where verification fails, the ticket is reopened and escalated rather than closed.

- **Example —** after a mailbox permission grant, the platform re-queries the target mailbox to confirm the permission is present before writing the closure note.

5.2 What flows alongside

Data	Written by	Used for
Ticket record	Ingestion and closure stages	Operational state, SLA measurement, reporting
Action trace	Execution stage	Evidence base for precedent retrieval and future planning
Evidence bundle	Decision core	Justifying each proposal to a reviewer and to audit
Audit record	Every stage	Compliance, forensic reconstruction, SIEM export
Feedback event	Approval and override actions	Retraining the model and tuning confidence thresholds

6 The Agent Orchestrator

The orchestrator is the control plane of the platform. It does not itself classify, plan or execute. It determines which specialist agent is required at each point in a ticket's life, invokes it, validates what comes back, maintains state, and enforces sequence. Where a ticket requires several systems to be touched in order, the orchestrator owns that ordering and the consequences of any step failing.

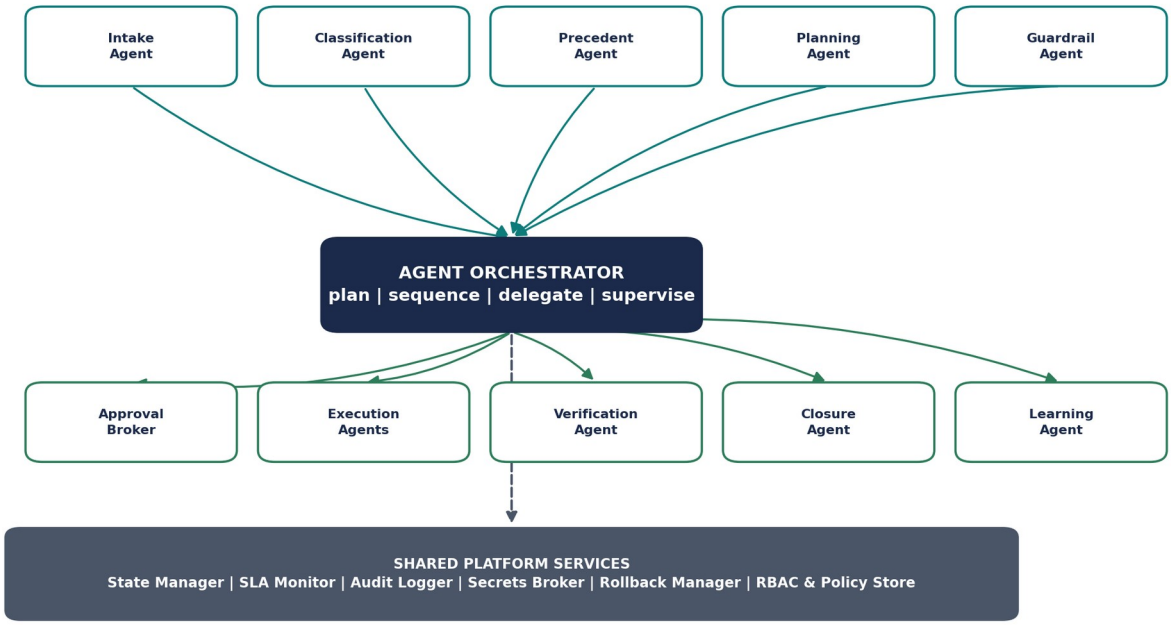

Figure 3 — Agent orchestrator and specialist agents

6.1 Why an orchestrator rather than a linear workflow

- **Multi-step work —** many tickets require several systems to be changed in a specific order, with verification between steps. A linear workflow cannot express dependency or recover cleanly from partial failure.
- **Long-running state —** a ticket may pause for hours awaiting approval or a change window. State must survive that wait, and the SLA clock must continue to run correctly.
- **Failure handling —** when a step fails at position four of six, the orchestrator determines whether to retry, roll back, or hand off, according to policy.
- **Concurrency —** independent steps execute concurrently where safe, and are serialised where not.
- **Traceability —** the orchestrator holds a single, continuous record of what was attempted, by which agent, with what result.

6.2 The specialist agents

Agent	Responsibility	Produces
Intake	Normalises the incoming record, resolves the requester and affected assets, detects duplicates	A clean, enriched ticket object
Classification	Determines category, subcategory, priority and owning tower	Taxonomy assignment with confidence
Precedent	Retrieves nearest historical resolutions and applicable knowledge articles	Ranked evidence set with success rates
Planning	Converts diagnosis into an ordered, executable action plan across systems	Step sequence with dependencies
Guardrail	Applies risk tiering, blast radius, deny-lists and approval policy	Routing decision — auto, approve or human
Approval broker	Presents the plan and its impact to the correct approver,	Approval, edit or rejection with reason

	captures the decision	
Execution	Domain-specific agents carrying steps to identity, network, cloud, endpoint, security and collaboration systems	Command results and system responses
Verification	Confirms the intended end state was actually reached	Pass, fail or partial verdict
Closure	Writes state and notes back to the system of record and notifies the requester	Closed ticket with full trail
Learning	Captures every human decision as a labelled training example	Feedback events for retraining

6.3 Shared platform services

Six services are available to every agent and are the reason behaviour remains consistent across them.

- **State manager —** holds the authoritative status of every in-flight ticket and guarantees state survives restarts and long waits.
- **SLA monitor —** tracks response and resolution targets, raises warnings before breach and drives escalation.
- **Audit logger —** writes the immutable record of every proposal, decision, command and response.
- **Secrets broker —** issues short-lived credentials to execution agents at the moment of use; credentials never appear in code, prompts or logs.
- **Rollback manager —** holds the validated reverse operation for every executable step and invokes it on failure.
- **RBAC and policy store —** resolves who may see, approve and execute what, mirroring Kyndryl's existing directory group structure.

6.4 Worked example — a firewall policy change

The following trace shows the orchestrator coordinating a genuine multi-system change from arrival to closure.

Step	Agent	Action
1	Intake	Ticket arrives requesting a geo-blocking rule for traffic from a specific region. Requester verified, affected estate resolved.
2	Classification	Classified to Security / Web Security / WAF Rule Change. Owning tower identified. Priority P2.
3	Precedent	Six comparable historical changes retrieved, together with the standard change article. All six used the same rule pattern.
4	Planning	Three-step plan produced — add rule at the perimeter firewall, update the cloud network security group, add the equivalent filter at the web application firewall.
5	Guardrail	Blast radius calculated at 4,200 users. Category flagged mandatory-approval. Rollback script generated and validated. Routed to APPROVE.
6	Approval broker	Tower lead is presented with all three commands, systems affected, users in scope, change window and rollback path. Approval recorded.
7	Execution	Steps executed in order through the firewall, cloud and WAF connectors. Each command and response logged.
8	Verification	Rule presence confirmed on all three systems. Test traffic from the blocked range confirmed rejected.
9	Closure	Ticket moved to Resolved with full notes. Requester and tower lead notified. Audit record sealed.
10	Learning	Approval captured with no edits — a positive example reinforcing this resolution path for this category.

7 Automation Model — Auto, Approve, Human

Every action the platform proposes is placed into one of three tiers. The tier is not a property of the ticket type alone; it is determined for each individual action, at the moment of decision, against current policy and current model confidence.

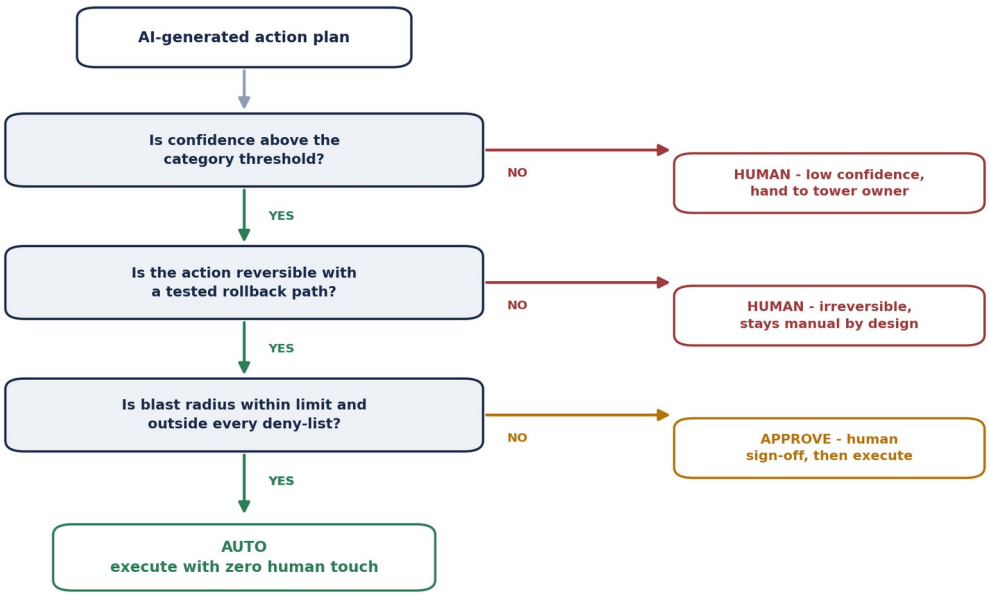

Figure 4 — How each action is routed to a tier

7.1 The three tiers

AUTO — zero human touch

The action executes immediately and the ticket closes without human involvement. To qualify, an action must clear three tests: model confidence above the threshold set for that category, a tested and available rollback path, and a blast radius inside the configured ceiling with no deny-list match.

- **Typical members —** password reset, account unlock, group membership addition, standard software installation from the approved catalogue, VPN profile push, mailbox and distribution-list access, disk cleanup, known-error remediation.

APPROVE — AI executes after human sign-off

The platform performs the analysis, retrieves the precedent, builds the complete plan and presents it for approval. The reviewer sees every command, the systems affected, the number of users in scope, the applicable change window and the rollback path before deciding. On approval the platform executes the plan; the reviewer performs no manual work. The reviewer may also edit the plan, and any edit is captured as a training signal.

- **Typical members —** new user onboarding across multiple platforms, firewall and WAF rule changes, geo-IP blocking, cloud resource provisioning, mailbox restore, non-standard software requests, application error remediation.

HUMAN — owned by people, by design

Certain work is excluded from automation as a matter of policy rather than capability. It is routed immediately to the correct named owner, with the platform's analysis and precedent attached so the engineer starts from an informed position rather than a blank ticket.

- **Typical members —** physical hardware dispatch and replacement, confirmed phishing and live security incident response, privileged and administrative access grants, data deletion and restore-over operations, server decommissioning.

7.2 How the tiers move over time

A category does not remain in a fixed tier. Every category begins in shadow, where the platform predicts but does not act, and its accuracy is measured against what engineers actually did. Once measured precision clears the promotion threshold, the category is promoted to assisted operation, and subsequently to automatic operation. Promotion is evidenced, governed and reversible; a category that degrades is demoted automatically.

Stage	Platform behaviour	Human involvement	Promotion condition
Shadow	Predicts and records; takes no action	Full — engineer works the ticket normally	Sustained precision above category threshold
Assisted	Builds the complete plan and presents it	Reviews and approves; performs no manual work	Sustained approval rate with low edit rate
Automatic	Executes and closes without prompting	Exception handling and periodic sampling only	Held under continuous monitoring

8 Ticket Catalogue by Domain

The following catalogue sets out the work the platform is intended to handle, with the tier proposed for each. Tier assignments are the recommended starting position and are confirmed against Kyndryl's own task-group data during the discovery stage of Milestone 1.

8.1 Access and Identity

The largest single domain by volume, and the source of the majority of early zero-touch coverage.

Ticket type	Tier
Password reset — directory and cloud identity	AUTO
Account unlock	AUTO
Multi-factor authentication reset and device re-registration	AUTO
Security group membership addition or removal	AUTO
Shared mailbox and distribution list access	AUTO
Collaboration workspace and document site access	AUTO
Application access grant against an existing entitlement	AUTO
New user onboarding — full access package across multiple platforms	APPROVE
User offboarding and access revocation	APPROVE
Role change — access realignment on internal transfer	APPROVE
Contractor account extension	APPROVE
Service account creation	APPROVE
Privileged access and credential vault requests	HUMAN
Local administrator rights grant	HUMAN

8.2 End User Computing

High volume, with a meaningful hardware component that remains human-owned.

Ticket type	Tier
Virtual desktop session failure and profile reset	AUTO
Printer and network drive mapping	AUTO
Mail client profile rebuild	AUTO
Disk space cleanup	AUTO
Device enrolment and compliance policy application	AUTO
Operating system update failure remediation	AUTO
Endpoint performance triage	APPROVE
Collaboration and productivity application failure	APPROVE
Device reimage	APPROVE
Crash and stability investigation	APPROVE
Hardware fault, replacement and physical dispatch	HUMAN
Peripheral provisioning and delivery	HUMAN

8.3 Software and Licensing

Highly deterministic; among the first categories promoted to zero-touch.

Ticket type	Tier
Standard software installation from the approved catalogue	AUTO
Licence assignment and reclamation	AUTO
Software removal	AUTO
Standard version upgrade	AUTO
Non-standard or unlisted software request	APPROVE
Application error and crash remediation	APPROVE
Application performance investigation	APPROVE
Licence renewal and true-up	APPROVE

8.4 Security Operations

Lower volume, higher consequence, and the domain where queue time carries risk.

Ticket type	Tier
Quarantined mail release against policy	AUTO
URL and domain block or unblock against an existing list	AUTO
Geo-IP blocking rule creation	APPROVE
Web application firewall rule change	APPROVE
Perimeter firewall policy request	APPROVE
Firewall-to-cloud policy push	APPROVE
Suspected phishing triage, containment and quarantine	APPROVE
Mail security sender allow and block	APPROVE
Endpoint detection alert triage and containment	APPROVE
Vulnerability remediation task	APPROVE
Certificate renewal and deployment	APPROVE
Confirmed phishing and live incident response	HUMAN
Credential vault and privileged session control	HUMAN
Data loss prevention policy exception	HUMAN

8.5 Network and Connectivity

Well-defined patterns with a clear boundary between routine and structural change.

Ticket type	Tier
Remote access client failure remediation	AUTO

Remote access profile push and reinstallation	AUTO
Wireless connectivity remediation	AUTO
Port opening request	APPROVE
DNS record creation and modification	APPROVE
Load balancer configuration change	APPROVE
Bandwidth and latency investigation	APPROVE
Site-to-site tunnel failure	HUMAN

8.6 Cloud and Infrastructure

Provisioning work with direct cost and security consequence, so weighted to the approve tier.

Ticket type	Tier
Virtual machine start, stop and restart	AUTO
On-demand snapshot and backup	AUTO
Tagging and cost-compliance remediation	AUTO
New virtual machine provisioning from an approved template	APPROVE
Virtual machine resize and scale	APPROVE
Storage expansion and disk addition	APPROVE
Resource group and subscription access	APPROVE
Restore from backup	APPROVE
Server decommissioning and data destruction	HUMAN

8.7 Collaboration and Messaging

Routine entitlement work with a small investigative tail.

Ticket type	Tier
Mailbox quota increase	AUTO
Calendar and mailbox delegate permission	AUTO
Meeting and recording access remediation	AUTO
Shared mailbox creation	APPROVE
Mail flow and delivery investigation	APPROVE
Mailbox restore	APPROVE

8.8 Monitoring and Alert-Driven Work

Machine-generated volume, handled without a ticket ever reaching a queue.

Ticket type	Tier
Capacity threshold alert — automated cleanup	AUTO
Service failure alert — automated restart with verification	AUTO
Alert triage and automated closure of confirmed false positives	AUTO
Certificate expiry alert remediation	APPROVE

9 Feature Catalogue

The solution comprises 100 features across twelve capability areas. Every feature is scheduled into one of three milestones according to its build complexity.

The complete feature-by-feature backlog — with description, complexity and scheduled delivery week for all 100 features — is provided as the accompanying workbook, AI Ops – Feature Plan.

9.1 Capability areas

The twelve capability areas below describe what the platform does. Section 10 sets out when each milestone delivers them.

Capability area	What it covers
Ticket Intelligence	Reading, understanding and classifying every incoming ticket, from any channel.
Knowledge & Precedent	Grounding each decision in Kyndryl's knowledge base and Kyndryl's own resolution history, with citable evidence.
Decision & Planning	Converting a diagnosis into an ordered, executable plan across one or more systems.
Guardrails & Risk	Constraining what the platform may do — risk tiers, blast radius, deny-lists and rollback.
Human-in-the-Loop	Everything a reviewer needs to approve, amend or reject an action with full context.
Execution	Carrying approved actions to the target systems and handling failure safely.
ITSM Operations	State, SLA, routing, ownership, escalation and write-back to the system of record.
Visibility & Reporting	Dashboards and reporting for engineers, tower leads and management.
Learning & Improvement	Turning every human decision into a training signal and governing the retrain cycle.
Security & Governance	Audit, access control, credential handling, residency and compliance evidence.
Platform & Deployment	Runtime, portability, scaling, resilience and programmatic access.
Advanced Operations	Multi-agent orchestration, problem and change management, and proactive remediation.

9.2 Distribution across milestones

The table below shows how the 100 features distribute across the three milestones, by capability area and by effort. *Note: This distribution and prioritization of features are subject to the final decisions by Kyndryl post the Due-Diligence Milestone,*

Capability area	*M1*	*M2*	*M3*	*Total features*
Ticket Intelligence	5	3	—	8
Knowledge & Precedent	5	1	1	7
Decision & Planning	1	3	2	6
Guardrails & Risk	3	4	1	8
Human-in-the-Loop	4	4	—	8
Execution	3	4	1	8
ITSM Operations	7	2	—	9
Visibility & Reporting	6	3	1	10
Learning & Improvement	1	5	1	7
Security & Governance	4	5	—	9
Platform & Deployment	3	3	1	7
Advanced Operations	—	—	8	8
Rollout & UAT	2	2	1	5
Total	44	39	17	100

9.3 How complexity was assigned

- **Easy** — capability that requires integration and configuration against a known pattern, with a well-understood interface and no dependency on model maturity.
- **Medium** — capability that spans several systems, requires policy evaluation or failure handling, or depends on a feedback cycle being in place.
- **Complex** — capability that coordinates multiple agents, reasons across tickets, or acts on signals before a ticket exists.

10 Delivery Milestones

Delivery is organised into three milestones of four weeks each, sequenced by build complexity rather than by business domain. The straightforward, high-volume capability is built first; medium-complexity governance and execution depth follows; advanced capability comes last. Every milestone closes with a production rollout and a user acceptance testing gate, so the platform is in live use from week four rather than at the end of the program.

Note: Delivery Milestone is subject to the findings of the Due Diligence Milestone.

Figure 5 — Three milestones of four weeks, sequenced by build complexity

10.1 Delivery cadence

Weeks	Milestone	Activity
1 – 3	Milestone 1	Build and configure the Milestone 1 feature set
4	Milestone 1	Production rollout, user acceptance testing and sign-off
5 – 7	Milestone 2	Build and configure the Milestone 2 feature set
8	Milestone 2	Production rollout, user acceptance testing and sign-off
9 – 11	Milestone 3	Build and configure the Milestone 3 feature set
12	Milestone 3	Production rollout, user acceptance testing and sign-off

Each rollout is cumulative. Milestone 2 is deployed on top of a Milestone 1 platform already running in production, and Milestone 3 on top of both. No milestone is a throwaway prototype.

10.2 Milestone 1 — Foundation and Core Automation

Weeks 1 - 4 | Complexity: Easy | 44 features

Establishes the platform, the full ITSM operations layer and the deterministic categories that already have a known resolution path.

Feature scope by capability area

Capability area	Features	Included in this milestone
Ticket Intelligence	5	Real-time ticket ingestion; Automatic classification; Duplicate and related-ticket detection; Requester and asset enrichment; Language and format normalisation
Knowledge & Precedent	5	Knowledge article matching with citation; Historical precedent retrieval; Resolution success scoring; Confidence scoring; Evidence bundle
Decision & Planning	1	Single-system action plan generation
Guardrails & Risk	3	Risk tiering; Deny-list enforcement; Mandatory-approval categories
Human-in-the-Loop	4	Approval workspace; Ownership and accountability display; Notify on completion; Full human override at any step
Execution	3	Core connector execution; Command-level logging; Real-time execution streaming
ITSM Operations	7	Full ticket state machine; Timestamped transitions with actor attribution; SLA engine; Assignment routing to owning queue; Named queue ownership; Write-back to the system of record; Requester notification
Visibility & Reporting	6	Overview dashboard; Live Ops Console; Ticket Explorer; Tower-lead dashboards; Category and subcategory surfaced throughout; SLA compliance reporting
Learning & Improvement	1	Shadow-mode evaluation
Security & Governance	4	Immutable audit log; Role-based access control; Secrets vault integration; Context redaction before model calls
Platform & Deployment	3	Containerised deployment; Multi-format data ingest; Automatic field mapping
Rollout & UAT	2	Tier A automation activation; Milestone 1 rollout and UAT

Week-by-week

	Activity	Features
Week 1	Build and configure — ITSM Operations, Knowledge & Precedent, Platform & Deployment, Security & Governance, Ticket Intelligence	10
Week 2	Build and configure — Decision & Planning, Execution, Guardrails & Risk, ITSM Operations, Knowledge & Precedent, Security & Governance, Ticket Intelligence, Visibility & Reporting	18
Week 3	Build and configure — Execution, Human-in-the-Loop, ITSM Operations, Knowledge & Precedent, Learning & Improvement, Visibility & Reporting	14
Week 4	Production rollout, user acceptance testing and sign-off	2

10.3 Milestone 2 — Assisted Execution and Governance

Weeks 5 - 8 | Complexity: Medium | 39 features

Adds the governance depth, multi-system execution and learning loop required for categories that need judgement.

Feature scope by capability area

Capability area	Features	Included in this milestone
Ticket Intelligence	3	Multi-channel intake; Urgency and impact detection; Missing-information detection
Knowledge & Precedent	1	Knowledge gap identification
Decision & Planning	3	Multi-system plan orchestration; Dependency resolution; Change window awareness
Guardrails & Risk	4	Blast-radius calculation; Per-category confidence thresholds; Rollback path generation and validation; Change-freeze enforcement
Human-in-the-Loop	4	Pre-approval impact panel; Plan editing; Rejection with structured reason; Delegation and escalation
Execution	4	Extended connector execution; Automatic retry with backoff; Failure handling and automatic rollback; Post-execution verification
ITSM Operations	2	Escalation management; Priority recalculation
Visibility & Reporting	3	Automation coverage reporting; Effort and savings reporting; Scheduled and exportable reports
Learning &	5	Feedback capture on every human decision; Scheduled retraining pipeline; Confidence

Improvement		threshold tuning; Precision and drift monitoring; Model versioning and rollback
Security & Governance	5	SIEM export; Short-lived credential issuance; Data retention and residency controls; Compliance and audit reporting; Segregation of duties enforcement
Platform & Deployment	3	High availability and horizontal scaling; API access; Disaster recovery
Rollout & UAT	2	Tier B assisted activation; Milestone 2 rollout and UAT

Week-by-week

	Activity	Features
Week 5	Build and configure — Decision & Planning, Execution, Guardrails & Risk, Ticket Intelligence	6
Week 6	Build and configure — Decision & Planning, Execution, Guardrails & Risk, Human-in-the-Loop, Learning & Improvement, Security & Governance, Ticket Intelligence	16
Week 7	Build and configure — Human-in-the-Loop, ITSM Operations, Knowledge & Precedent, Learning & Improvement, Platform & Deployment, Security & Governance, Visibility & Reporting	15
Week 8	Production rollout, user acceptance testing and sign-off	2

10.4 Milestone 3 — Advanced Operations and Autonomy

Weeks 9 - 12 | Complexity: Complex | 17 features

Extends the platform from resolving tickets to reducing the number of tickets raised, and into problem and change management.

Feature scope by capability area

Capability area	Features	Included in this milestone
Knowledge & Precedent	1	Automated knowledge authoring
Decision & Planning	2	Alternative plan proposal; Cross-ticket correlation
Guardrails & Risk	1	Automatic demotion on degradation
Execution	1	Concurrent execution where safe
Visibility & Reporting	1	Executive summary view
Learning & Improvement	1	New category discovery
Platform & Deployment	1	Multi-tenant operation
Advanced Operations	8	Multi-agent orchestration for complex tickets; Problem management support; Change management integration; Proactive and predictive remediation; Conversational intake; Capacity and cost optimisation; Continuous compliance monitoring; Long-tail category mining
Rollout & UAT	1	Milestone 3 rollout and UAT

Week-by-week

	Activity	Features
Week 9	Build and configure — Advanced Operations, Decision & Planning	4
Week 10	Build and configure — Advanced Operations, Decision & Planning, Execution	5
Week 11	Build and configure — Advanced Operations, Guardrails & Risk, Knowledge & Precedent, Learning & Improvement, Platform & Deployment, Visibility & Reporting	7
Week 12	Production rollout, user acceptance testing and sign-off	1

10.5 Rollout and user acceptance testing

Each milestone concludes with the same four-part gate. Sign-off at this gate is the condition for starting the next milestone.

5. **Production rollout** — the milestone's features are deployed into the production environment, alongside everything already live from earlier milestones.
6. **User acceptance testing** — Kyndryl's nominated users exercise the delivered features against agreed acceptance criteria, using real tickets in the live environment.
7. **Defect resolution** — any defects raised during UAT are corrected and retested within the gate week.
8. **Sign-off** — formal acceptance of the milestone, confirming the feature set is delivered and operating as specified.

Acceptance criteria for each milestone are agreed in writing before that milestone begins, so UAT tests against a defined standard rather than an open-ended review.

10.6 What remains human throughout

Across all three milestones, a defined population of work is never automated. Physical hardware dispatch, confirmed phishing and live security incident response, privileged and administrative access grants, data deletion and restore-over operations, and server decommissioning remain human-owned. This is a policy position and is not revisited by model maturity.

11 Continuous Learning

The platform improves through use. Every interaction between an engineer and a proposal is a labelled example, and the accumulation of those examples is what allows automation coverage to rise without further development work.

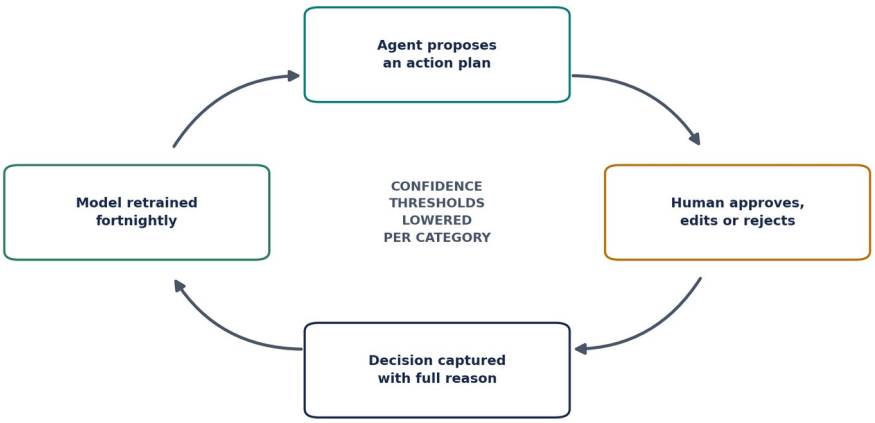

Figure 6 — The continuous learning loop

11.1 How the loop operates

9. **Proposal —** the agent produces a complete plan with its supporting evidence and a confidence score.
10. **Human decision —** an engineer approves it, amends it, or rejects it with a structured reason.
11. **Capture —** the decision is recorded together with the full context that produced it — a clean approval is a positive example, an amendment identifies precisely where the plan was wrong, a rejection with reason identifies why.
12. **Retraining —** accumulated feedback is incorporated on a fortnightly cadence, with every model version retained and reversible.
13. **Threshold adjustment —** as measured precision rises within a category, the confidence threshold for that category is lowered, moving a greater share of its volume from assisted to automatic operation.

11.2 Why amendments matter more than approvals

A clean approval confirms an existing path. An amendment is more valuable: it identifies the exact point at which the platform's understanding diverged from an experienced engineer's judgement, and provides the correction. Amendment rate is therefore tracked as a primary quality measure, and a falling amendment rate within a category is the principal evidence supporting promotion of that category to a higher tier of autonomy.

11.3 Governance of the loop

- **Versioning —** model versions are retained and any version can be reinstated.
- **Evaluation gate —** each retrained model is evaluated against a held-out set before promotion.
- **Drift monitoring —** precision is monitored continuously per category, with alerting on degradation.
- **Automatic demotion —** a category whose precision falls below threshold reverts automatically to assisted operation.
- **Change record —** threshold changes are recorded, attributable and reversible.

12 Integration Landscape

Connectors carry approved actions to the target system using that system's own API or command interface. Each is a discrete, independently deployable component, so scope can be extended or re-pointed without disturbing the platform.

Domain	Systems	Representative actions	Milestone
Identity	Entra ID, Active Directory, Okta	Password reset, account unlock, group membership, MFA reset, entitlement grant and revocation	**M1**
Endpoint & VDI	Intune, SCCM, Citrix, VMware Horizon, Azure Virtual Desktop	Device enrolment, policy application, session reset, profile repair, software deployment	M1/M2
Cloud	Microsoft Azure (ARM), Azure Front Door WAF	Resource provisioning, resize, storage expansion, network security group change, snapshot	M1/M2
Network & Security	Palo Alto, Fortinet, Cloudflare WAF, F5 BIG-IP, Microsoft Sentinel	Firewall rule change, geo-IP blocking, WAF policy, load balancer configuration, DNS change	**M2**
Collaboration	Exchange Online and Microsoft 365	Mailbox permission, distribution list, quota, shared mailbox creation, mail flow remediation	M1/M2
SIEM & Endpoint Security	Splunk, CrowdStrike, Microsoft Defender,	Alert triage, host containment, evidence collection, backup and restore operations	**M2**

	Veeam		
Service Management	ServiceNow	Ticket read, state update, work notes, resolution write-back, assignment and escalation	**M1**

Additional connectors are added as scoped extensions. Because every connector implements the same internal contract, adding a system is a bounded integration task and does not require change to the decision, guardrail or orchestration layers.

13 Security, Governance & Compliance

13.1 Control position

Control	Position
Audit logging	Every proposed and executed action written to an immutable log recording actor, system, command, evidence and rollback path. Queryable in the platform and exportable to Kyndryl's SIEM.
Access control	Role-based access mirroring Kyndryl's existing service management and directory group structure. The platform never holds broader privilege than the human role whose work it performs.
Policy enforcement	A guardrail engine sits between every proposal and any execution. Deny-lists, blast-radius ceilings and mandatory-approval categories are absolute and are not overridden by model confidence.
Credential handling	All connector credentials held in a secrets vault and issued short-lived at point of use. Credentials never appear in code, prompts, logs or configuration files.
Data residency	The ticket store, action traces, audit log, classification model and precedent index are held entirely inside Kyndryl's environment. The reasoning step calls a hosted GPT API with minimised, redacted context and no retention for model training; it can be replaced by a self-hosted model where policy requires.
Redaction before transmission	Identifiers, credentials and free-text personal data are stripped from any context sent to the reasoning API. Redaction rules are configurable and every outbound call is logged.
Segregation of duties	Proposer, approver and executor are distinct roles and cannot be collapsed by configuration.
Human override	Any action may be stopped, amended or reversed at any point. Every override is logged and returned to the learning loop.
Retention	Retention periods for tickets, traces and audit records are configurable to Kyndryl's regulatory position.

13.2 Operating assurance

- **Evidence before autonomy —** no category is granted autonomy without measured evidence of precision in shadow and assisted operation.
- **Reversibility —** every executable step carries a validated reverse operation before it is offered for approval.
- **Fail safe —** connector failure, model unavailability or policy ambiguity results in a hand-off to a named human owner, never in an unlogged or partial action.
- **Change control —** all thresholds, deny-lists and approval categories are recorded, attributable and reversible.

14 Infrastructure and Deployment Model

The platform is built, deployed and operated on Kyndryl's own infrastructure. Compute, storage and network capacity are provisioned inside Kyndryl's cloud account, and remain under Kyndryl's ownership and administrative control throughout.

14.1 Deployment principle

- **Kyndryl's data stays with Kendryl—** the ticket store, action traces, audit log, classification model and precedent index are held inside Kyndryl's infrastructure boundary at all times.
- **The reasoning model is consumed as a service —** the reasoning step calls a hosted GPT API. Only minimised, redacted context is sent, and it is not retained for model training. No GPU or model-hosting capacity is required.
- Kyndryl **holds the account —** capacity is provisioned in accounts Kyndryl own, so it can be re-purposed, re-sized or decommissioned at Kyndryl's discretion.
- Kyndry **retains what is built —** the deployed platform, its configuration and the trained models are Kyndryls, and continue to operate independently of any ongoing relationship.

14.2 Environments

Environment	Purpose	Availability
Development	Build, integration testing, connector validation, configuration and UAT	Single instance; can be scaled down or suspended outside build and test windows
Production	Live ticket processing, execution against connected systems, audit and reporting	Continuously running

14.3 Indicative sizing

The platform is deliberately lean. Because no model is hosted locally, there is no GPU node and no high-memory inference tier. Final sizing is confirmed during discovery.

Component	Specification	Production	Development
Application and orchestration nodes	4 vCPU, 16 GB RAM	2 nodes	1 node
Database — PostgreSQL with vector extension	4 vCPU, 16 GB RAM, 250 GB SSD	1 instance	1 small instance
Cache and message queue	2 vCPU, 8 GB RAM	1 node	shared
Object storage — evidence, exports, backups	Object storage	500 GB	100 GB
Monitoring and log aggregation	Managed / lightweight	Included	Shared
Load balancer, egress and backup	Managed services	Included	Included

14.4 Where it runs

Any provider below runs the same containerised platform without code change. The choice is a cost and policy decision, and is reversible later.

Provider	Position	Best suited to
DigitalOcean / OVH	Recommended	Sufficient for this workload at roughly half the cost of a hyperscaler; simplest commercial terms
Azure / AWS	Alternative	Where an enterprise agreement, existing landing zone or a procurement policy mandates a hyperscaler

On-premise	Where mandated	Where a policy requires physical control of the hardware; costed separately on request

15 Estimated Infrastructure Cost

Infrastructure is procured by Kyndryl, in Kyndryl's own account, and billed directly to Kyndryl. It does not pass through ThirdEye Data and carries no markup. Costs are shown per month, for the two environments plus the GPT API, against two provider options.

15.1 Development environment — per month

Component	Azure / AWS	DigitalOcean / OVH
Application node	₹14,000	₹6,000
Database	₹12,000	₹5,000
Cache and message queue	₹4,000	₹1,500
Object storage	₹500	₹300
Network, backup and monitoring	₹3,500	₹1,200
Development subtotal	₹34,000	₹14,000

15.2 Production environment — per month

Component	Azure / AWS	DigitalOcean / OVH
Application and orchestration nodes	₹28,000	₹12,000
Database	₹32,000	₹14,000
Cache and message queue	₹8,000	₹3,500
Object storage	₹1,500	₹500
Load balancer, egress and backup	₹8,500	₹3,000
Monitoring and log aggregation	₹6,000	₹3,000
Production subtotal	₹84,000	₹36,000

15.3 GPT API — per month

The reasoning model is consumed through the GPT API on the lowest-cost model tier, which is sufficient for classification support, planning and summarisation at this volume. Cost is usage-based and scales with ticket volume rather than being a fixed capacity charge.

Basis	Position
Tickets processed per month	~16,700
Average model calls per ticket	4 (assist, plan, summarise, verify)
Average tokens per call	~3,000 input + ~600 output
Model tier	Lowest-cost GPT tier (mini class)
Computed monthly cost	~USD 55
Provisioned with headroom for retries and peaks	₹10,000 per month

15.4 Total monthly cost

	Azure / AWS	DigitalOcean / OVH
Development environment	₹34,000	₹14,000
Production environment	₹84,000	₹36,000
GPT API	₹10,000	₹10,000
Total per month	₹1,28,000	₹60,000

Total per year	₹15.4 L	₹7.2 L
Total over three years	₹46.1 L	₹21.6 L

15.5 Recommendation

DigitalOcean is the recommended option. It runs the same platform on the same containers, meets the performance requirement at this volume, and costs ₹68,000 per month less than an equivalent hyperscaler build — ₹8.2 L a year, or ₹24.5 L over three years.

Because the platform is containerised, this is not a one-way decision. If an enterprise agreement, a landing-zone standard or a procurement policy later requires Azure or AWS, the same build moves across without code change, and the cost position simply shifts to the first column.

15.6 Basis of estimate

- Sizing supports a ticket volume of approximately 200,000 per year, with headroom for growth.
- No GPU or model-hosting capacity is required. The reasoning model is consumed as a hosted GPT API and billed on usage.
- Costs are indicative list-price estimates in Indian Rupees and exclude taxes and duties.
- Infrastructure is provisioned in Kyndryl's own account and billed directly to Kyndryl. It does not pass through ThirdEye Data and carries no markup.
- The development environment can be scaled down or suspended outside build and testing windows, reducing its cost further.
- GPT API cost is usage-based and moves with ticket volume. Current published rates should be confirmed at contracting, as provider pricing changes periodically.
- Final sizing is confirmed during discovery, once connector scope and data retention requirements are fixed.

16 Engagement and Commercial Model

The engagement is a development and customisation programme. ThirdEye Data designs, builds, configures and deploys the platform on Kyndryl's infrastructure; Kyndryl provides the environment it runs in and meets the infrastructure cost directly.

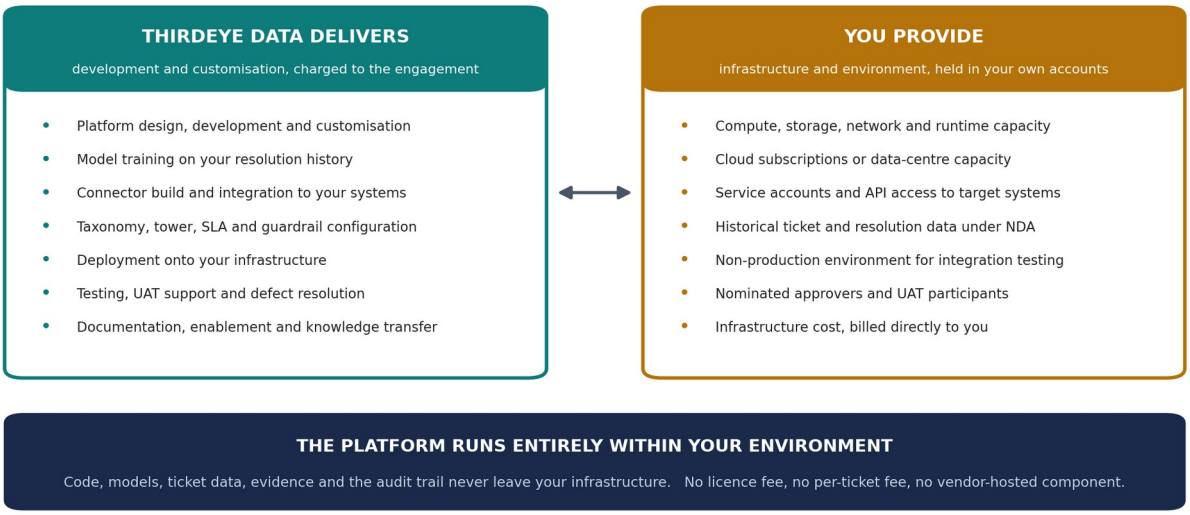

Figure 7 — Responsibility and cost split across the engagement

16.1 What ThirdEye Data delivers

Service	Description
Discovery and assessment	Analysis of ticket history and task groups to establish the automatable population, tier assignment and baseline effort position.
Data onboarding and model training	Ingestion of historical ticket and resolution data under NDA, and training of the classification and precedent models on Kyndryl's environment.
Platform development and customisation	Build and configuration of the feature set defined in Section 9, tailored to Kyndryl's taxonomy, towers, policy and approval structure.
System integration	Connector development and authentication against Kyndryl's instances, with testing in the development environment before production use.
Deployment	Installation, configuration and commissioning of the platform on Kyndryl's infrastructure across both environments.
Shadow operation and validation	Operation across the full ticket flow without acting, to establish measured precision per category before promotion.
Phased activation	Category-by-category promotion to assisted and then automatic operation, on evidence.
Testing and UAT support	Test execution, UAT facilitation and defect resolution within each milestone gate.
Knowledge transfer	Documentation, administrator enablement and operational handover.
Post-go-live support	Optional — ongoing support, retraining, new category onboarding and enhancement, contracted separately.

16.2 What Kyndryl provide

- **Infrastructure —** compute, storage and network capacity across the development and production environments, provisioned in Kyndryl's own account.
- **Environment and cost —** cloud subscriptions or data-centre capacity, and the associated infrastructure cost, billed directly to Kyndryl.
- **Access —** service accounts and API access with appropriate privilege on the service management platform and each in-scope target system.

- **Data —** three months of ticket exports and resolution action history under NDA for model training.
- **People —** nominated approvers per category and per tower, and participants for user acceptance testing at each milestone gate.

16.3 Charging basis

Cost element	Borne by	Basis
Development and customisation	Charged by ThirdEye Data	Fixed charge per milestone, invoiced on UAT sign-off at the end of each four-week milestone
Infrastructure	Paid directly by Kyndryl	Actual consumption in Kyndryl's own accounts, at Kyndryl's negotiated rates, with no markup
Post-go-live support and enhancement	Charged by ThirdEye Data	Optional, contracted separately after Milestone 3 acceptance
Licence or subscription fee	Not applicable	There is no licence fee, subscription fee or per-ticket charge

Detailed development and customisation charges are set out in the accompanying commercial annexure.

16.4 Payment linked to acceptance

Development charges are structured against the three milestones described in Section 10. Each milestone concludes with a production rollout and user acceptance testing, and the associated invoice falls due on written sign-off of that milestone. Acceptance criteria are agreed before each milestone begins, so payment is tied to demonstrated delivery rather than elapsed time.

Milestone	Weeks	Payment trigger
Milestone 1	1 – 4	UAT sign-off at the end of week 4
Milestone 2	5 – 8	UAT sign-off at the end of week 8
Milestone 3	9 – 12	UAT sign-off at the end of week 12

16.5 Ownership

On acceptance of each milestone, the platform as deployed in Kyndryl's environment — the configuration, the trained models, the connector integrations and the data they operate on — is Kyndryl'ss. It runs on infrastructure Kyndryl own and continues to operate independently of any ongoing commercial relationship. There is no dependency on a hosted service, a licence key or a recurring fee for the platform to remain in operation.

17 Assumptions and Dependencies

17.1 Assumptions

- Ticket volumes, task-group distribution and handling times stated in this document are indicative and are confirmed against Kyndryl's data during discovery.
- Three months of historical ticket data with associated resolution and action detail is available for model training.
- Tier assignments in Section 8 are a recommended starting position and are confirmed jointly before activation.
- Feature scheduling across the three milestones is indicative and is confirmed during discovery.

- Infrastructure estimates in Section 15 are indicative list-price figures and exclude taxes, duties and any discounts already available to Kyndryl's organisation.
- Work classified as human-owned in Section 7 is excluded from automation coverage.

17.2 Dependencies on Kyndryl's organisation

- Infrastructure provisioned and available in both environments before the relevant milestone begins.
- API access to the service management platform with sufficient privilege for read and write-back.
- Service accounts and appropriate privilege on each in-scope target system.
- Confirmation of the tower, queue and ownership structure to which tickets are routed.
- Nominated approvers per category and per tower for the assisted tier.
- A development environment available for connector testing prior to production activation.
- Executed non-disclosure agreement covering the transfer of historical ticket data.
- Nominated business and technical owners, UAT participants, and an agreed governance cadence.

17.3 Out of scope

- Procurement, ownership and ongoing cost of infrastructure, which is met directly by Kyndryl's organisation.
- Replacement or re-platforming of the existing service management platform.
- Remediation of underlying infrastructure faults that generate ticket volume.
- Physical hardware supply, logistics and field services.
- Changes to the target systems themselves beyond the actions the platform is configured to perform.

18 Next Steps

The following steps convert this document into a firm, committed delivery plan with a validated baseline.

	Step	Purpose
1	Confirm connector scope	Identify the specific service management instance, identity provider, network and security platforms, cloud environment and VDI systems in scope for each milestone.
2	Share task-group detail	Provide the activities L1 and L2 perform per task group, so that handling times and tier assignments are validated against Kyndryl's actual operation rather than assumed.
3	Confirm ownership structure	Confirm the tower, queue and named-owner structure to which tickets are routed, and nominate approvers per category.
4	Confirm infrastructure position	Select the deployment option, confirm the target environment, and validate the sizing in Section 14 against Kyndryl's existing capacity and any discounts already in place.
5	Provide historical data	Three months of ticket exports and resolution action history under NDA, so that models and the effort baseline are built on Kyndryl's real environment.
6	Baseline and lock the plan	Re-run the coverage model on validated data and issue a committed milestone plan with agreed acceptance criteria and payment schedule.

Once steps 1 to 5 are complete, a committed delivery plan with locked milestone dates, agreed acceptance criteria and a confirmed infrastructure position is issued for signature.