# Agentic Journeys in Citizen Services

*A policy note on trusted **Agentic Public Services***

**Towards Building a Trust Infrastructure at Scale**

**Prepared by Ashish Srivastava and Amit Prakash**

**Working Draft | For Discussion**

International Institute of Information Technology Bangalore · 2026

# Contents

Executive Summary ..... 3

1 The Citizen Burden ..... 4

2 The Anatomy of an Agentic Journey ..... 7

3 Agentic Journeys in Motion ..... 9

4 What an Agentic Journey Does That a Workflow Cannot ..... 12

5 An Architecture That Holds an Agentic Journey ..... 14

6 Trust and Assurance Across the Journey ..... 17

7 The Human Dividend ..... 20

8 Towards Shared Journey Infrastructure ..... 22

9 An Invitation ..... 23

Annex A: Six journeys used to test the argument across domains ..... 24

References ..... 25

About IIIT-Bangalore and the Authors ..... 27

# Executive Summary

## The citizen burden

India's digital public infrastructure has made individual steps faster and more dependable: identity is verified in seconds, benefits reach bank accounts directly, and documents travel with the citizen. The difficulty now lies between those steps. A school admission, a child's immunisation, a pension or an insurance claim may be digital at every point and still leave the citizen to repeat information, follow up across offices, interpret silence and absorb each exception herself.

Agentic AI, which can plan and act towards a goal rather than only answer a prompt, is improving quickly at short tasks. Citizen journeys are harder. They run for months, pause for long periods, change direction and are full of exceptions, while the citizen waits for an outcome rather than a record of completed steps.

## The journey should be the unit of trust

This note proposes that the journey, rather than each transaction or agent, become the unit of design, governance and trust. A governed journey stays with the citizen over time, coordinates across institutions and reasons within limits it cannot change for itself. Each step is rule-based (a set rule applies), autonomous (the system works out the step within set boundaries) or human (a person decides). The clean path is mostly rule-based; what carries real value or risk is mostly autonomous or human. Holding a journey this way needs a few clear responsibilities, described in Section 5, rather than a new central system.

## Governance is an important part of design

Models will keep improving on their own schedule driven by a global industry. The layer that decides who may act, whose approval counts, how consent lasts for months and who notices when a journey stalls is a matter of public trust, and public institutions are best placed to shape it. Seen through the journey, consent stays meaningful for as long as it is needed, authority is bounded and visible, risk is reassessed as the journey progresses, people are placed at consequential moments, ~~and~~ failure announces ~~itself.~~itself, and what the system infers about a citizen is treated with the same care as what she provides. Assurance then attaches to the whole journey, not only to each agent (Section 6).

## The human layer is an essential part of governance

None of this diminishes people. Agentic systems can handle hard cases too, but the authority to make a decision binding, and the judgement of whether it is fair and safe, stay with a person. The same responsibilities recur across sectors: reviewing an agent's reasoning, handling exceptions, supporting citizens through difficult moments and checking that journeys keep their promises. Together they form a new layer of skilled public-service roles, worth building with intent (Section 7).

## Shared services at state or central level

A citizen's protection should not depend on which programme serves her. Rebuilding consent, delegation, traceability and recourse for every service is costly and leaves safeguards uneven. Building them once, as shared services at state or central level held by a neutral, publicly accountable steward, lets every programme inherit safeguards already tested elsewhere, while ministries and states keep their own journeys and decisions (Section 8). India has done this before with identity and data exchange; the journey between a citizen and her outcome may be the natural candidate next.

# 1 The Citizen Burden

*Public systems are increasingly good at completing transactions. Citizens still carry the burden of completing the journey.*

A parent receives a message saying that her child has been selected under an education entitlement. The school says verification is still pending. The income certificate carries one spelling of the family name, the birth record another. A deadline is approaching, the portal offers no explanation, and every office can see only its own step. The parent has not encountered a failed transaction. She has encountered a journey that nobody is holding.

This is a familiar pattern across public services. Digital systems have become increasingly capable at identity checks, applications, payments, document exchange and status display. New agentic tools can now perform short chains of work, often impressively: interpret an intent, call a service, compare options, prepare a form and complete a task within one session.

Many public services have this journey-shaped character. The six examples below show how widely the pattern can occur.

🏫 **Education Inclusion**: A low-income parent applies for a school place under the Right to Education framework. Between the application, the lottery, document verification and the school's confirmation, she must track each stage herself and correct any mismatched documents. If the school delays or refuses, she has to work out on her own where to turn.

💉 **Universal Immunisation**: A caregiver must keep a child on a vaccination schedule that runs through the early years of life. If the family moves for seasonal work, a session is missed or two districts hold different records, it falls to her to notice, find a new session site and explain the child's history all over again.

🤰 **Maternal Health**: A woman with a high-risk pregnancy moves between antenatal visits, tests, frontline workers and, sometimes, referral facilities. No one may be watching the whole picture. A danger sign between visits, a missed check-up or a referral that was never completed can go unnoticed until it becomes an emergency.

🩺 **Frontline Worker Enablement**: An ASHA (Accredited Social Health Activist) serves many households across immunisation, pregnancy and other programmes, each with its own register, application and reporting cycle. She becomes the human link between them, piecing together priorities and follow-ups from memory and paper, with little help in seeing which family needs her most this week.

🏦 **Dignity in Pension**: An elderly beneficiary relies on a monthly pension that depends on annual life verification and a correctly linked bank account. When a payment stops after a small record mismatch, no one tells her why. She must discover it herself, travel to the bank or district office, and explain her situation again at each counter.

🌱 **Resilient Rural Livelihoods**: A farmer whose crop is damaged must file a claim within tight time limits, gather evidence, wait for an assessment and match land records held by different offices. If the figures do not align, the claim stalls or is disputed, and he is left to chase several institutions, often far from home, to learn what went wrong.

These six examples are illustrative rather than exhaustive. Many other citizen-impacting journeys unfold through hospitals, financial institutions, government departments, schools, universities and other public or regulated institutions. The principles discussed here apply equally wherever citizens must navigate a long-running service across multiple steps, actors or systems.

The note draws selectively on the six examples because they are familiar and reveal the coordination burden in an ordinary citizen's life. Each brings a different aspect of the journey into view, but none limits the wider relevance of the approach.

## From transactions to journeys

To understand why these services need a journey lens, it helps to distinguish how public systems organise delivery from how citizens experience it. Systems often divide a service into separate transactions and tasks. A citizen, however, experiences a single continuous path from a need to the outcome she is trying to secure.

Digital support for citizen services can be understood at three levels:

- **Transaction.** A point event: authenticate, submit, pay or retrieve.
- **Task agent.** A short chain of work, usually within a single session or a bounded authority.
- **Agentic journey.** A path that may unfold over weeks or months, cross institutions and channels, pause for long periods and change direction when circumstances change. Through all of this, it remains connected to an outcome that no single step can deliver.

The technology world has become very good at transactions and is moving quickly on short agentic tasks. Long journeys remain underdeveloped. Even sophisticated demonstrations tend to compress complexity into one session: the user asks, the agent acts, and an answer arrives. The harder question is whether a governed system can still hold the citizen's outcome after four dormant weeks, a change of district, a missing document, a disputed decision, and a handoff to an official working in another system.

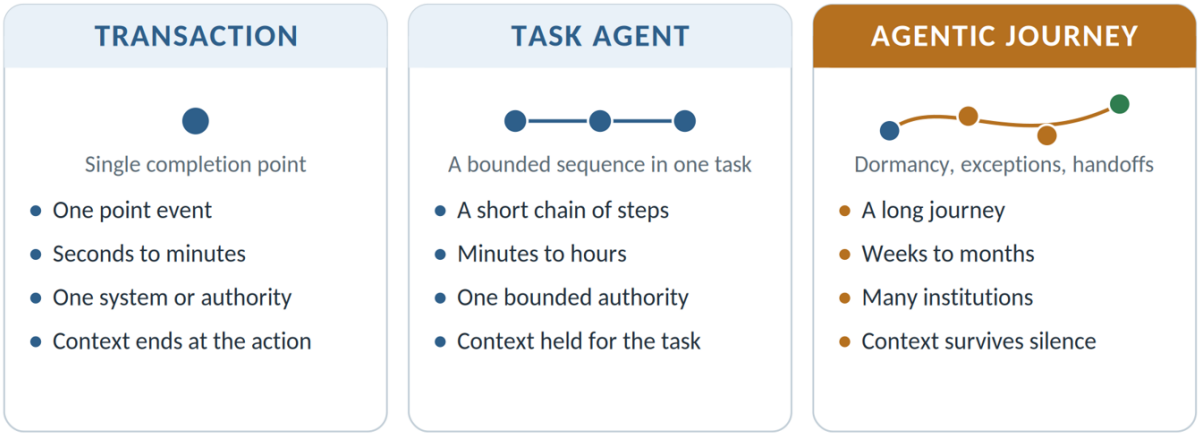



***Figure 1**. From transactions to journeys: transaction, task agent and agentic journey.*

## Why citizen journeys are the hardest case

Citizen journeys are the hardest version of this problem, and it is worth seeing why. Long journeys exist within enterprises too, in procurement cycles or insurance claims, but an enterprise usually sits under a single authority that can redesign the process end-to-end. Citizen journeys have no such luxury: they run across the longest durations, carry the most fragmented authority, face the widest range of exceptions, and fall on the people least able to absorb coordination failure (NITI Aayog, 2025).

The citizen, therefore, becomes the integration layer. She remembers which office asked for what, carries context across portals, chases status, interprets silence, and begins again when an exception

breaks the expected path. Frontline workers often carry the same burden on her behalf, stitching together registers, calls, messages and informal knowledge.

## Why this space remains open

At the time of writing, three structural reasons may help explain why long-running journeys remain an underappreciated frontier.

- **First**, journey competence is domain competence. Immunisation requires the grammar of schedules, minimum intervals, contraindications, session calendars and field practice. An insurance claim requires a different grammar of evidence, verification, assessment and dispute. This depth does not generalise easily to a single product.
- **Second**, authority is fragmented. No single buyer or department owns the whole outcome, so the market receives no clear demand signal for it.
- **Third**, the economics favour tools more readily than sustained outcomes. It is easier to sell a capability into each department than to remain answerable for a citizen's result across several of them.

None of these is simply a gap that newer, better generative AI models can fill. Better models help, but they do not decide who may act, whose approval carries legal or administrative weight, how consent survives over time, or who notices when the journey has silently stopped. That is why the question belongs as much in the digital public infrastructure conversation as in the AI product conversation (RIS, 2024).

There is also a question of timing. Agentic systems are already appearing in public-service settings, and they will continue to do so as the technology matures. The real choice is not whether this happens, but whether the institutional side of it is shaped with intent or assembled piece by piece by whoever arrives first. Shaping it deliberately is what keeps a citizen's protection even, rather than leaving it to the accidents of procurement.

Before the space can be occupied deliberately, the object itself needs a clearer definition. The next section defines the agentic journey and its key terms in simple, policy-friendly language.

# 2 The Anatomy of an Agentic Journey

## What is a citizen journey?

Before considering what makes a journey agentic, it helps to identify what makes it a citizen journey. A citizen journey is the complete path from a need to an outcome when that path spans time, institutions, or moments of uncertainty. The citizen experiences it as a single engagement, even when government systems divide it into multiple transactions.

Not every service is a journey. A certificate download should remain a download, and nothing here argues for agentic treatment of what is genuinely transactional. The lens applies where the properties below hold; applying it elsewhere would add machinery without adding value.

A journey is not merely a long transaction. Six properties distinguish it, and together they form the vocabulary this note uses to describe citizen journeys: Section 3 traces them through three journeys, and Section 6 derives the governance obligations from them.

| Property of a citizen journey | What it asks of a well-designed journey | Example from citizen services |
|---|---|---|
| **A journey is long, and mostly dormant.** | Hold the context through long gaps, and resume without asking the citizen to begin again. | A pregnancy moves between scheduled visits, tests and long intervals at home. |
| **The next steps depend on what happens along the way.** | Adapt the next step to new facts, without inventing authority. | An RTE admission changes after the lottery, verification, a document mismatch or a school refusal. |
| **Exceptions and surprises are inherent in every journey.** | Treat exceptions as part of the service, since this is where most value and risk sit. | A crop claim becomes a dispute when evidence, assessment and land records do not align. |
| **Many institutions share the journey, each owning only its own steps.** | Coordinate across offices, and keep track of who may decide what, so that the outcome has an owner. | A frontline health worker, such as an ASHA in India, prepares an update to the record; an authorised official attests it. |
| **The stakes can change midway.** | Reassess risk as facts change, and adjust oversight and the human role accordingly. | A routine pregnancy is reclassified as high-risk in the sixth month. |
| **The citizen is assured of an outcome, not a set of completed steps.** | Track the outcome, and notice when it fails to arrive, including when the failure is silent. | A pension stops after a bank-seeding mismatch, although no official decided to stop it. |

***Table 1.*** *Six properties of a citizen journey, and what each asks of a well-designed journey.*

## When does a citizen journey become agentic?

The six properties explain why citizen journeys strain today's systems. Each asks for something that transaction systems and fixed workflows were not built to do: hold context through long silences, adapt as facts change, treat exceptions as normal work, coordinate across many hands, adjust oversight as the stakes rise, and notice when an outcome fails to arrive. Today, these demands fall largely on the citizen and on frontline workers, who carry the journey between offices.

A citizen journey becomes agentic when the system begins to meet these demands itself. It holds the goal, preserves the relevant context, works out what should happen next and coordinates the capabilities needed to move towards the outcome. It does so under human authority, and within limits that it cannot change on its own.

The capability that makes this possible is an agent. Put simply, an agent is AI that can be given a goal and work towards it, deciding the steps along the way within limits set for it (UK Government, 2025). In a citizen journey, one agent usually stays with the citizen throughout: the companion agent. It keeps the thread of the journey and calls on other, more specialised capabilities when they are needed.

Five terms recur in the rest of the note. The last three describe how a step is determined and where authority sits; they do not imply that autonomous action is ungoverned.

- **Agent**: An AI-enabled system that can understand a goal, plan the steps needed, and use available information, tools or services to pursue it within defined limits.
- **Companion agent**: An agent that stays connected to the citizen throughout the journey. It preserves relevant context, tracks progress and coordinates the capabilities needed to reach the outcome.
- **Rule-based action**: An action determined by an established rule or instruction, sometimes described as a configured or deterministic step. The agent does not need to work out a new path for the particular situation.
- **Autonomous action**: A decision or proposed action that the agent works out for the particular situation, without a person directing each step. It must remain within rules, permissions and boundaries established in advance.
- **Human action**: A validation, approval, decision or conversation that requires a person's authority, judgement or empathy.

# 3 Agentic Journeys in Motion

Agentic journeys differ in more than length. Each combines the six properties from Section 2 in its own way, and so each brings its own mix of complexity, stakes and rhythm. Some journeys are tested by surprises along the way, such as a family that migrates partway through a vaccination schedule. Some are tested by stakes that rise without warning, when a routine case becomes urgent. Others are tested by time, running quietly for years until something expected fails to happen.

This section follows three journeys that make these differences visible. Child immunisation is described in detail, from birth to completion. A high-risk pregnancy and a continuing pension are summarised more briefly, to show how the same approach adapts when the stakes or the rhythm change.

## Journey 1. Child immunisation: managing real-world surprises

Immunisation offers a particularly clear view of the difference between a workflow and a journey. It combines hard-governed facts, such as the national schedule and minimum intervals between doses (MoHFW, 2023), with life circumstances that no rulebook can enumerate, such as migration, missed sessions, changing availability and caregiver hesitancy. The persona below is illustrative; the pattern is not.

Meena's daughter is born at a district facility. Registration captures the child's details and Meena's consent: what information may be used, for what purpose and for how long. A companion agent establishes the immunisation plan from the governed schedule and the local session calendar. Before the six-week session, a reminder reaches Meena through the channel she chose, such as WhatsApp, SMS or a voice call, using text or recorded voice according to her preference. The visit happens, and the record is updated through the health worker's attestation. The schedule, session sites and consent record are Rule-based. The official update remains Human because Meena cannot write directly to the official record.

### Dormancy, then resumption

Then, for a month, nothing happens. That is the point. When the ten-week dose approaches, the companion resumes as if no time had passed. It knows what was given, what is due, which channel Meena prefers and what she asked about last time. A family function leads to a missed dose, and the companion offers the next permissible session. This remains Rule-based because the situation fits a known recovery rule.

### Migration: reasoning through an exception

The family then migrates for seasonal work, first to another district and later across a state boundary, with two doses pending. No configuration entry describes this child, this dose history, this destination and this week's session availability. The system needs to reason through four questions:

- Which doses can be caught up, and in what sequence, without violating minimum intervals?
- How can the record follow the child?
- Which nearby session site has capacity?
- Do the accumulating misses indicate vulnerability?

A catch-up plan is generated, checked against the governed clinical constraints and surfaced to the receiving health worker. She validates it against what she knows of local logistics and what the family can realistically manage. The reasoning is Autonomous; the validation is Human. The schedule could be configured. The migration could only be reasoned about.

### The moment that should remain human

After a local rumour about an adverse event, Meena stops responding. The system notices a pattern of silence: unanswered reminders, a missed session and no request to reschedule. It does not respond by increasing the volume of automated messages. It routes the case, with its history, to the health worker, because a hesitancy conversation belongs to a person rather than a prompt. A supervisory medical officer separately reviews the adverse-event flag.

### Closer attention, completed journey

Repeated misses move the journey into a flagged tier, bringing closer human attention. With the health worker's support, the catch-up sequence is completed. The journey closes not because a form was submitted, but because the child's immunisation outcome is current. Along the way, the system coordinated with scheduling, health-record and vaccine-readiness systems that it did not own or control.

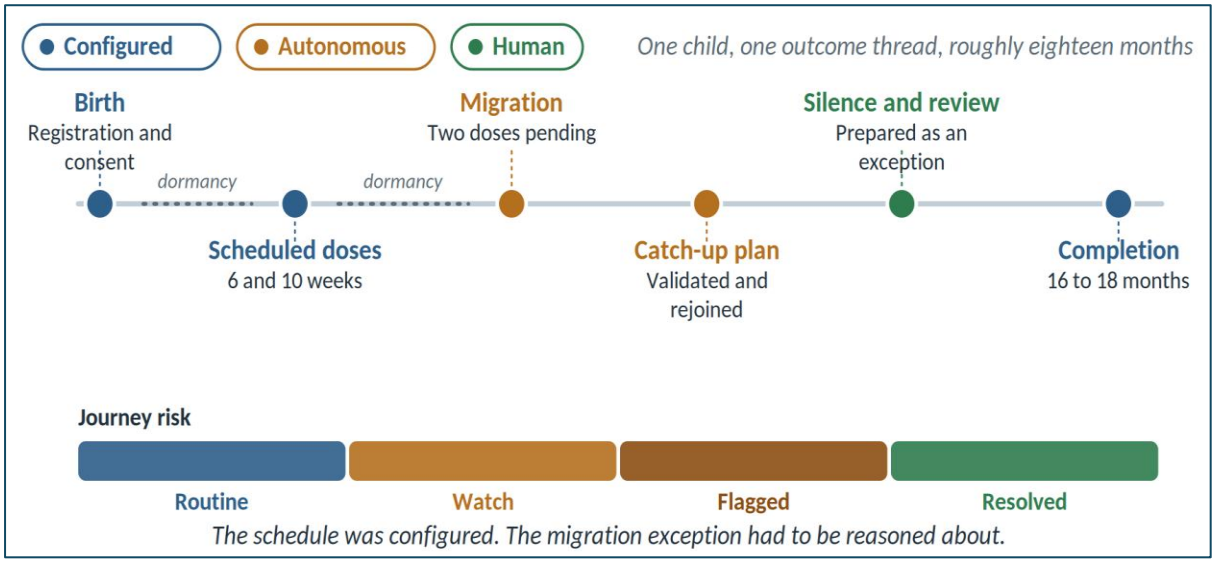



***Figure 2**. The immunisation journey across Rule-based, Autonomous and Human modes.*

### Summarised steps of the immunisation journey

| Step | What happens | Mode |
|---|---|---|
| 1 | Registration captures child details, channel preferences and journey-scoped consent. | **Rule-based** |
| 2 | The plan is drawn from the governed immunisation schedule and local session calendar. | **Rule-based** |
| 3 | A missed routine session is moved to the next permissible slot under a stored recovery rule. | **Rule-based** |
| 4 | Migration with two doses pending triggers a context-specific catch-up plan within minimum-interval constraints. | **Autonomous** |
| 5 | The receiving health worker validates the plan against the family situation and local feasibility. | **Human** |
| 6 | A pattern of silence after an adverse-event rumour is detected and prepared as an exception case. | **Autonomous** |
| 7 | The health worker holds the hesitancy conversation; a medical officer reviews the adverse-event flag. | **Human** |
| 8 | Repeated misses shift the journey into a flagged tier, tightening oversight and changing the human role. | **Autonomous** |

| Step | What happens | Mode |
|---|---|---|
| 9 | The official record is updated through authorised attestation and the next milestone is established. | **Human** |
| 10 | The journey closes when the governed outcome is complete, with its decision and consent trail retained. | **Rule-based** |

***Table 2.*** *The clean path is largely Rule-based. The moments that create value and risk are predominantly Autonomous or Human.*

*Although specialised capabilities may handle registration, scheduling, notification, readiness checks and record synchronisation, the citizen should not have to navigate that organisational chart. She should experience one relationship that remains connected to the outcome.*

## Journey 2. High-risk pregnancy: when the stakes rise midway

Sunita is in her second pregnancy in a rural block. At first her care is routine: antenatal visits drawn from the national guidelines, reminders by voice call in her own language on a basic phone, and a record kept alongside the systems health workers already use (MoHFW, 2024). Across a high-risk pregnancy, visits, tests, home follow-ups and unplanned episodes can commonly reach 20 to 40 or more touchpoints, spread across several facilities and frontline workers (an illustrative range). Today, Sunita and her family are left to hold them together.

In the sixth month, a danger sign is reported between visits. The system recognises that the risk picture has changed and proposes that the pregnancy be reclassified. A medical officer confirms it, and the safeguards tighten with the stakes: more checkpoints, shorter escalation times and a named person responsible. The system prepares the referral, identifies the appropriate facility, checks that it is ready to receive her and tracks the ambulance request. The medical officer at the referral unit decides where she goes.

Later, a high-risk visit is missed with no request to reschedule. The silence is noticed, and the case is assembled with its history, so that the health worker begins from the full picture. When Sunita hesitates over a referral, that conversation is held by a person, not a prompt. The journey closes when she has delivered safely and post-delivery follow-up is complete, not when a form is submitted. A routine pregnancy reclassified in the sixth month is the same journey with different stakes, and its governance has to change with it.

## Journey 3. Pension continuity: lifting the coordination burden from the citizen

An elderly widow's pension journey has an annual heartbeat: life verification, eligibility continuity and twelve expected payments. Its characteristic failure is not a dramatic error but an absence. A payment silently stops after a bank-seeding mismatch. No component reports failure because every component is technically available. Today, she is usually the first to notice. She must then visit the bank, the district office or a service centre, explain her situation from the beginning at each, and keep following up until someone resolves it.

The companion's most valuable act is to notice that something expected did not happen, assemble the grievance with the journey history attached and route it to the accountable officer, so that she no longer has to carry the case herself.

The immunisation journey alone has displayed all six properties from Section 2. Notice where the modes fell: the clean path was mostly Rule-based; nearly everything that mattered was Autonomous or Human. The two shorter journeys show the same pattern under different pressures: in the pregnancy, the governance tightened as the stakes rose; in the pension, the companion lifted from the citizen a burden she would otherwise have carried alone. The natural question is whether a sufficiently elaborate workflow could do the same. The next section addresses that question directly.

# 4 What an Agentic Journey Does That a Workflow Cannot

A capable workflow team, looking at Meena's journey, might reasonably propose adding a migration rule. That rule would help. The difficulty appears in what follows: a family that migrates twice, a contraindication that arises during catch-up, a session cancelled after the plan is compressed, a record present in one state and missing in another, or a caregiver who can attend only on a day that conflicts with the technically optimal sequence. Each situation can be configured individually. The question is what happens to a system that tries to configure them all.

## The rulebook keeps growing

The same trap is easiest to count outside citizen journeys, in a platform that prepares classroom material for teachers. Consider four subjects across five grades and sixteen output types. That already creates 320 configuration entries for one state board before classroom content is authored. A second language medium takes the count beyond 600. A second state, with a different pedagogy, calendar and assessment structure, doubles it. Textbook editions, local variations and district calendars push the count into the thousands.

A system that handles diversity through exhaustive rule-authoring gradually stops behaving like a software product and becomes a permanent editorial operation. This is not only a cost problem. It is a structural trap. The education and immunisation examples are different domains, but they expose the same limit: every new circumstance creates another branch to author, validate and maintain.

## Some knowledge exists only when the situation arises

Even an unlimited configuration budget would not close the gap, because some of the knowledge a situation requires does not exist until the situation exists. Meena's catch-up plan depended on this child's precise dose history, this family's destination, that district's current availability and the constraints that applied at that moment. A rule-writer would have needed to know the future.

What can be written in advance are the constraints the plan must respect: minimum intervals, contraindications, authority boundaries, required validations, and escalation conditions. What cannot be written in advance is the plan itself. This distinction between knowledge held as governed fact and judgement generated when context arrives is the core case for agentic systems in long journeys.

## Configure the boundaries, not the paths

Bounded autonomy turns the configuration problem inside out. It allows a system to reason within limits it cannot cross, with defined moments where a person decides. Instead of attempting to configure every possible path, the design configures the boundaries: mandatory rules, action limits, quality gates, escalation triggers and human approvals. These are comparatively stable, inspectable and reviewable. New situations can then be handled through reasoning within the same governed boundaries rather than through endless new branches.

> **A PRINCIPLE FOR GOVERNMENT**
>
> **Govern the authority limits, not every scenario**
>
> *No one can write rules for every situation a citizen may face. What government can set, and check, are the limits of authority: what the system may never do, which facts it must respect and when a person must decide.*

## Why the timing has changed

The idea of managing a long case, rather than a single transaction, is not new. Case-management practice recognised decades ago that the long-running case was the right unit of analysis. What it

lacked was a way to act on that insight at scale. Its limitation was not conceptual. Without a capable reasoning layer, every unmodelled situation, even if simple, was handled by fixed rules or still returned to a person for resolution. The coordination burden never truly moved away from the human actors. The case file followed the journey, but a human still carried it.

What has changed is the reasoning layer. Systems that can now reason within explicit constraints make the journey question newly answerable. That is also what makes it pressing, because the same capability will be applied to public services regardless of whether the institutional design has kept pace. The governance questions that once felt academic have become immediate and practical.

Something else has changed alongside the reasoning. Until recently, systems struggled to keep the thread of a long, intermittent engagement and lost the story across gaps in time. Newer models, and better ways of helping a system remember and resume its work, are easing that limit (Gemini Team, 2024; Anthropic, 2025a; Du, 2026). A companion that remembers what was done and picks up weeks later, without asking the citizen to start again, is becoming practical rather than aspirational.

If bounded reasoning is to carry the journey, something needs to hold the context, hold the rules, enforce the boundaries and preserve the points where a person decides. That is an architecture question.

# 5 An Architecture That Holds an Agentic Journey

An agentic journey needs somewhere to live and something to hold it. The architecture is described here by its responsibilities, not by any particular model, protocol or product, because those will change while the responsibilities will not.

Departments already run systems of record built around their own responsibilities, and India's digital public infrastructure connects identity, payments and data exchange between them; what remains largely unconnected is the journey itself. Nothing in this pattern replaces existing portals, official records or departmental authority. It coordinates what already exists. Experience from building and operating AI platforms at population scale, together with analysis of the six citizen journeys, suggests five elements.

### 1. A companion that stays with the citizen

The companion is the citizen's continuing point of contact. It understands her in her own language, holds the goal across months, picks up after long gaps without asking her to retell the story, and recognises when a routine matter has become an exception. It coordinates and prepares; it does not decide. Nothing it proposes becomes official merely because it proposed it.

### 2. A workspace where the journey lives

Each journey needs a light workspace that holds its current state: what is done, what is pending, which approval is waiting and which exception is open. It is where a caregiver confirms, a health worker validates and an officer approves, and it is the memory from which the system resumes when the journey stirs again.

The workspace is not a parallel register. Official records stay in the systems that own them, which preserves institutional accountability while making fragmented authority workable.

### 3. A foundation of fixed rules

Every domain has settled knowledge that the system must treat as fact rather than reason about: an immunisation schedule and its minimum intervals, an eligibility rule, a map of who may approve what. This layer applies that knowledge; it does not invent it. Its value lies in being traceable to an authoritative source and kept current with the institutions that own it.

Journey competence is therefore domain competence. A journey is only as good as the depth and currency of the domain knowledge beneath it, not the fluency of a general-purpose agent.

### 4. Reasoning that works within the rules

The agentic layer does the work that cannot be specified in advance: building a catch-up plan after migration, assembling a grievance from a long history, or replanning when circumstances change. It reasons on the rules, not around them.

Reasoning about an action and carrying it out are two separate events, and a check sits between them. Before any consequential action, the system confirms that it is permitted, that consent still holds and whether a person must approve it. In Meena's case, the catch-up plan was checked against the clinical intervals and sent to the receiving health worker; the agent could not update the official record itself. This check turns a boundary into something enforceable rather than advice that the agent may ignore.

The boundaries themselves take four forms:

- **What the system may never do,** such as write to an official record without attestation.
- **What it may never produce,** such as a plan that breaks clinical intervals.

- **When it must stop and hand over to a person,** such as on repeated misses, vulnerability or a dispute.
- **How much it may do at once,** so that automated activity never outruns what the human layer can absorb.

## 5. Governance that runs through everything

Governance establishes who may act, what limits apply, where a person must decide and how the journey can be examined or challenged. It runs through every element rather than sitting beside them. Two capabilities matter most.

### Traceability: every consequential action can be reconstructed

Traceability means that any consequential action can be reconstructed long after it happened. For such an action, the system should be able to say what was done, on what basis, under which version of the rules, with whose consent and on whose approval, even months later (Wang et al., 2026; EU AI Act, 2024).

A concrete case shows what this is worth. Consider a disputed crop-insurance settlement: the trace can show the claim as filed, the evidence assembled, the mismatch detected, the verification ordered, the officer's decision with reasons, and the exact update sent to the insurer. The dispute then becomes a review of a documented decision path rather than an archaeological dig across four offices.

Transaction systems log events. Journeys need the thread that connects events across time and institutions, including the moments where humans approved or overrode a proposal. That is a stronger discipline than conventional logging.

### Journey observability: are citizens progressing?

Observability is the ability to see how a journey is going at any moment. In practice, it means monitoring whether citizens are actually progressing towards their outcomes: which cases have stalled, which are looping between the same steps, and where silence has replaced the expected next step. A district where claims repeatedly stall at one verification stage then becomes visible as a policy signal, not only an operational one.

Together, traceability and observability make risk control demonstrable, compliance verifiable and audit affordable. Without them, oversight would depend on trusting the system's good behaviour, which is not a durable foundation for public deployment.

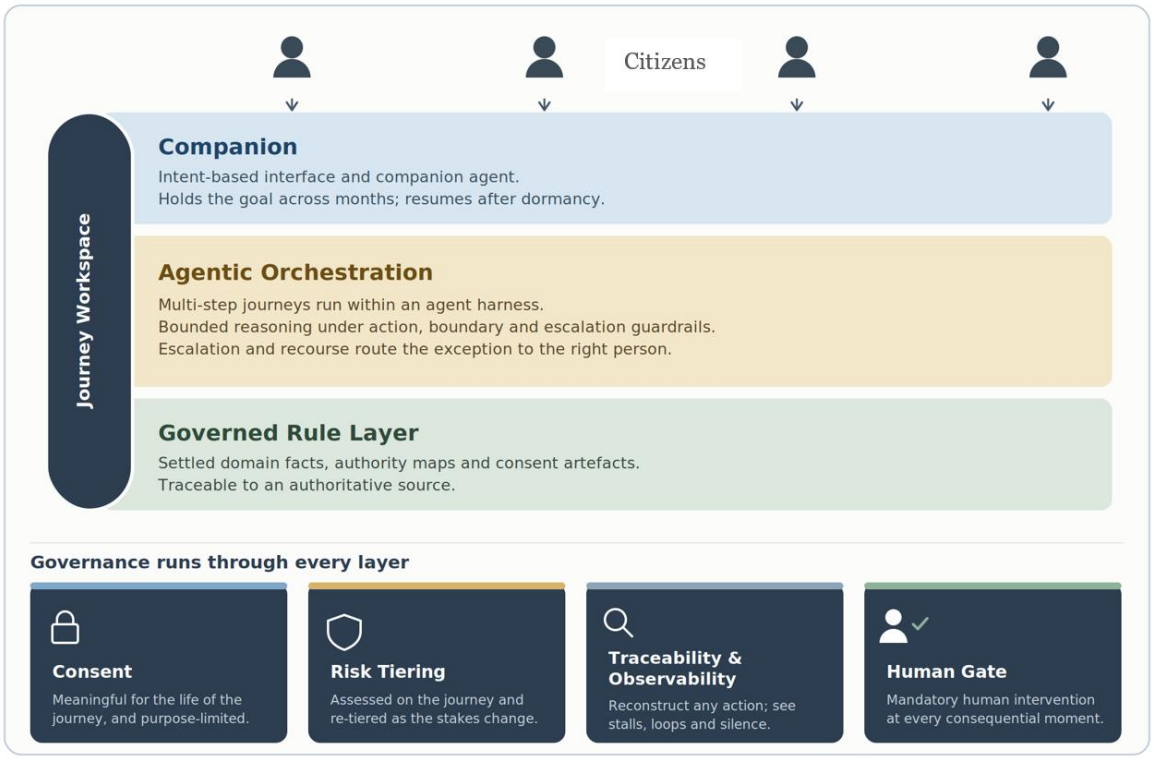



***Figure 3**. Reference architecture for a governed agentic journey.*

## Rules and reasoning run through every layer

The most important idea here is easy to miss. The line between what is fixed, or deterministic (rule-based), and what is reasoned does not separate one system from another; it runs through each of them. The immunisation schedule is fixed, while the catch-up plan built on it is reasoned. The review gate is fixed, while the plan it examines is reasoned. The architecture makes this line visible at every layer, so that anyone reviewing the system can see where a rule ends and reasoning begins.

None of this requires a new portal, a new official register or, necessarily, a central platform.

---

*Holding a journey, however, is not the same as making it trustworthy. The next question is what the journey owes the citizen.*

---

# 6 Trust and Assurance Across the Journey

The governance question can begin with a general checklist applied to each component, or with the properties of the journey and what each one requires. This note takes the second route, because governance derived from the journey travels more naturally across schemes and states, while a component checklist may leave the citizen's outcome unexamined.

### 1. Consent lives for the full duration of the journey

Consent for a journey cannot work the way consent for a single transaction does. A form signed once at a counter may authorise one action and then be spent. A journey may act on a citizen's behalf for months. Its consent must therefore remain specific to purpose, verifiable when consequential actions arise, easy to review or withdraw, and capable of being reaffirmed when the journey changes.

The consent Meena gave at birth may cover scheduling, reminders and record linking. If repeated missed doses move the journey into closer vulnerability follow-up, that change in purpose and intensity should be visible to her rather than hidden inside the original grant. This extends familiar consent-artefact thinking by adding duration and revisitability (DPDP Act, 2023; DPDP Rules, 2025). It also offers a practical way to reconcile two legitimate goals: continuity requires remembering, while data minimisation requires forgetting. Journey-scoped consent ties retention to the journey's life and purpose. What a journey remembers can also include what it infers. When a companion concludes that a family may be vulnerable, as in Meena's case, that assessment is not a verified fact. It should be handled with the same care as the information the citizen provides: used only for the purpose authorised, visible to her, and open to challenge before it changes how she is treated.

### 2. Authority is delegated, bounded and visible

As a journey crosses roles and institutions, the person or agent doing the work is often not the authority that can make it official. That gap is normal, but it must remain visible. Every consequential step should be able to answer: on whose behalf did this happen, what exactly was permitted, and was that authority still valid at the time? A system may prepare an update to an official record, but execution should pass through the approval or attestation of the authorised role, with the mandate and approval preserved as part of the decision trail.

A frontline worker may know that a record needs correction and may prepare the update with supporting context. The authorised officer's attestation carries it into the official record, while the trace preserves both contributions. Public-service journeys therefore need more than agent identity.

**Capability is not authority.** An agent may be able to find a record, prepare a correction or call an institutional service, but that does not give it permission to act. Permission comes from an explicit mandate, setting the purpose, duration, permitted actions and the points where a person must approve, which the agent cannot enlarge for itself. As authority passes between agents or institutions it should narrow rather than expand, and its withdrawal should reach every link in the chain.

### 3. Risk is reassessed as the journey progresses

Classifying the model or agent is not sufficient. The same system can carry journeys of very different consequences, and a single journey can change consequences in flight. A routine pregnancy and a high-risk pregnancy may use the same companion, but the appropriate human checkpoints, escalation speed and review depth are not the same.

The more useful unit is therefore the journey instance, supported by a risk tier and defined triggers for re-tiering (NIST, 2023; NIST, 2024). When the stakes rise, safeguards tighten automatically: closer human review, shorter escalation windows, a narrower scope for autonomous reasoning and a clearer line of accountability.

### 4. Human judgement is placed where it matters

Human-in-the-loop is easy to claim and easy to hollow out. A person clicking approve hundreds of routine items is oversight in name but not in substance (EU AI Act, 2024). A governed journey places people at moments of consequence: approval into official records, vulnerability flags, contested claims, ambiguous eligibility, caregiver hesitancy and grievances.

The quality of the placement matters as much as its location. The person should receive a prepared case: the journey history, evidence, relevant provisions, unresolved mismatch and decision points isolated. The decision remains genuinely hers. Placement is a design question; the people who fill these positions are considered in the next section.

### 5. Failures are caught before the citizen finds them

The citizen should not be responsible for discovering that her journey has failed. A journey owes her visible failure: stalls detected when nothing happens on schedule, loops detected when she is sent around the same step, and recourse that begins from the journey's own record rather than from zero at a counter.

The silently stopped pension payment becomes a raised exception and an assembled grievance. A claim that circles between verification and resubmission becomes a visible loop. A school admission waiting beyond the verification window becomes an ageing case on an accountable queue. Outcome accountability is therefore not only a matter of who made a wrong decision. It is also the ability to notice that no decision has been reached.

This obligation matters most for the citizens least able to absorb a broken journey. The people whose journeys leave the expected path are more often migrant families, older citizens, and households with irregular records or limited digital access. They stand to gain the most from a journey that holds itself together, and they are the first to be harmed when it fails silently. Visible failure is therefore a question of fairness as much as of efficiency.

## Assurance covers the whole journey

Because these obligations belong to the journey as a whole, assurance should attach to the whole journey and not only to each agent. A review can examine who is accountable, what outcome is intended, how consent and authority are delegated, where people decide and how a citizen can seek recourse. It should be proportionate to risk: a review before deployment, monitoring during operation, and a fresh look after a material change or a serious incident.

A light registry of agents can record who owns each one and what it may do, while a governing authority approves the journey's boundaries and requires stronger controls for higher-risk journeys. A journey reviewed once in this way, such as immunisation, could then be adapted by several states around a shared governed core.

Periodic review needs to be matched by checks at the moment of action, which decide whether each consequential step stays within the approved design (Schmitz et al., 2025). Each needs the other: review alone is only a statement of intent, and checks alone never ask whether the rules are still right. Encouragingly, the tools that coordinate an agent's work are making such checks and human gates easier to build and to inspect (Anthropic, 2024; Anthropic, 2025b).

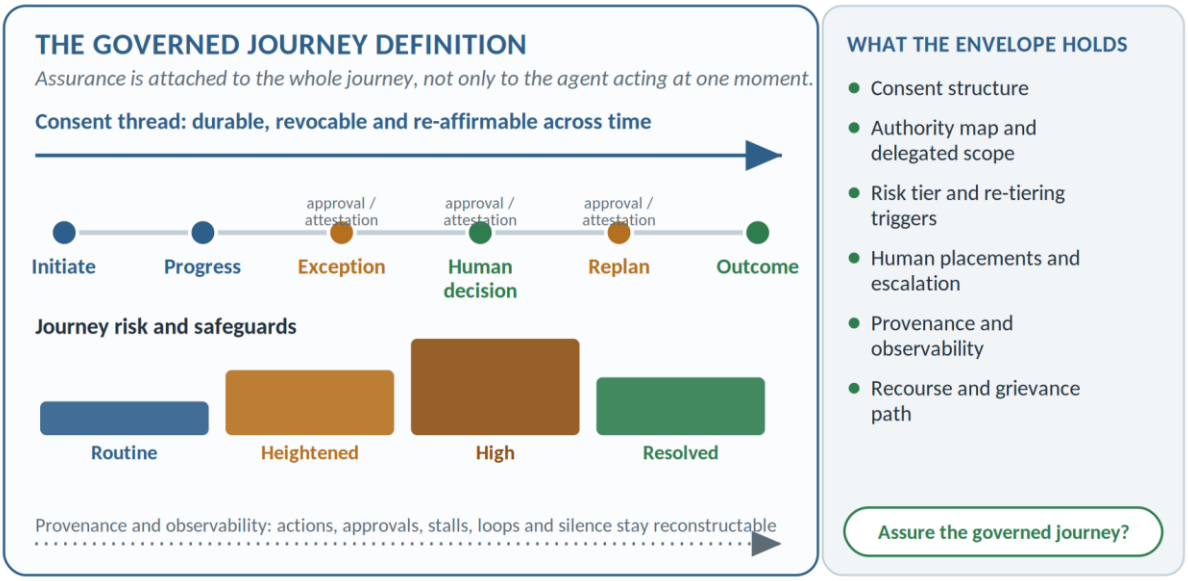



*Figure 4. The governed journey as a possible unit of assurance.*

## How does this complement existing AI governance?

India's own approach to AI governance offers a natural point of connection here. The India AI Governance Guidelines set out a principle-based framework, built on seven guiding principles, the sutras, and organised under six governance pillars (MeitY, 2025). They place trust, accountability and a people-first orientation at the centre, and they are well suited to questions of models, actors and individual actions. Those foundations remain essential.

The journey lens can be read as a way of carrying those same principles across time, into the space between transactions where consent must stay meaningful, authority is delegated across institutions, risk can change midway and failure may appear as silence. Understood this way, the obligations above complement the wider framework rather than compete with it, and the idea of assuring a whole governed journey sits comfortably alongside the framework's own emphasis on assurance.

Every promise in this section ultimately ends with a person: someone approves, receives an escalation, has a difficult conversation, reviews a dispute or answers the citizen. The human layer is not an afterthought to the architecture. It is part of the governance itself.

# 7 The Human Dividend

Because the human layer is part of the governance rather than a supplement to it, it is worth looking closely at what that layer actually contains. Every governed agentic journey holds decision points where a person's authority, judgement or empathy is required. The system should recognise these moments, prepare the relevant context, and bring the matter to someone equipped and authorised to act. Human involvement, on this understanding, is not a temporary arrangement that better technology will eventually retire.

It would be a mistake to read this as machines taking the simple work while people keep the difficult work. Agentic systems are increasingly capable of the hard parts, too. They can reason through a tangled case, coordinate across several institutions and arrive at a defensible course of action faster than any person could. The distinction that matters is therefore not difficulty but responsibility. What stays with a person is the authority to make a decision binding, the duty to answer for it, and the judgement of whether an outcome is fair, safe and within the rules a public service is expected to honour. A system can work out what is possible, and often what is likely correct. A person remains accountable for whether it is acceptable.

In practice, this places people at a recognisable set of moments. Someone gives an update its official force and stands behind it. Someone settles a contested claim and owns that decision, whatever the agent has proposed. Someone holds the conversation with a hesitant caregiver, where trust matters more than throughput. And someone stays answerable when a journey touches a vulnerable citizen or departs from what was expected. The common thread is not that these things lie beyond a machine. It is that responsibility for them should rest with a person.

There is a second consequence that deserves equal attention. Because these moments recur across journeys and across sectors, they add up to a coherent body of skilled work rather than a scattering of occasional tasks. Reviewing an agent's reasoning, owning a structured exception, supporting a vulnerable citizen, and checking that a journey has kept its promises are all capabilities that can be developed, supported with the right tools, and given a settled institutional home. This layer of roles is therefore something to build with intent, and the rest of this section describes the shapes it tends to take.

> THE HUMAN LAYER
>
> **Human roles keep autonomy bounded**
>
> *Governed agentic systems move human effort from producing the outcome to taking responsibility for it, and that responsibility remains permanent because the governance remains permanent.*

Across the journeys examined, the human work settles into four recurring roles. How a state staffs them, through existing personnel, redesigned responsibilities or new positions, remains a matter for institutional choice.

## The person who helps the citizen through

Some citizens will always need a person to help them through, whether from limited digital access, low confidence, or a journey that has begun to slip. That is the work of a Citizen Support Facilitator, based at common service centres, panchayat offices, district helpdesks or assisted service points. The role supports citizens with limited digital access or confidence, helps with document readiness and guided steps, and reaches out when a journey shows someone slipping. For many citizens, the most trusted interface to a digital companion will remain a person. Assisted access is therefore a first-class channel, not a fallback.

### The person who owns the exception

An Escalation Coordinator sits at a block or district desk and takes end-to-end responsibility for cases the system hands over: repeated rejections, vulnerability flags, disputes and cross-department knots. The role does not override policy. It makes sure the case reaches the authority that can decide, with its context intact, and that the decision returns to the citizen rather than disappearing into another queue.

### The person who keeps the system honest

An Agent Support Specialist, often drawn from existing management information system (MIS) or programme teams, watches journey observability: whether citizens are progressing, not only whether systems are running. The role notices the district where claims stall at one step, a data feed drifting out of sync, or an exception pattern that indicates a systemic issue. These observations feed back into rules, thresholds, integration quality and programme design.

### The person who checks that journeys keep their promises

A Policy Compliance Officer works across oversight functions, reviewing decision trails, sampling high-consequence outcomes, watching for exclusion or bias, and verifying that the obligations described in Section 6 exist in practice rather than only in design documents.

### Technology multiplies each person's reach

These roles are themselves AI-enabled. An escalation coordinator does not receive a raw complaint. She receives a structured exception packet containing the journey history, evidence, relevant provisions and decision points. Her workspace consolidates cases and deadlines. A compliance officer does not reconstruct a file from paper and memory; the journey trace brings the decision path to the screen in minutes.

---

*"Agents orchestrate the journey; people remain accountable for the decision."*

---

The same technology that creates the need for deliberate human judgement also multiplies each person's reach. That is what makes the human layer a plausible public function rather than a staffing fantasy. It also dissolves the usual binary between AI taking jobs and AI creating jobs. The more defensible observation is that governed agentic systems move human effort from carrying out the work to taking responsibility for whether it is right.

This shift asks for new capabilities: reading agentic behaviour, handling structured exceptions, consent and data-protection fluency, and the judgement to recognise when the case packet is not enough. Training, tooling and progression deserve the same design attention as the technology (McKinsey, 2025; CSET, 2024; IndiaAI, 2026).

The roles, obligations and architecture recur across health, education, social protection and livelihoods. That repetition raises a practical policy question: how much should each programme build for itself, and how much might usefully be shared?

# 8 Towards Shared Journey Infrastructure

Every governed journey depends on the same protections. Building them once, as shared services at state or central level, is the surest way to keep a citizen's protection consistent and to let each new programme start from safeguards that have already been tested. This is a strategic choice for those who oversee several programmes or steward digital public infrastructure.

The case for sharing comes from recurrence. Immunisation, maternal health, pensions and crop insurance serve different purposes, yet each needs consent that stays meaningful, authority that stays within its limits, decisions that can be explained and failures that do not stay hidden.

Rebuilding these protections separately for every service is costly. More importantly, uneven safeguards create uneven rights, because the quality of a citizen's protection would then depend on which programme happens to serve her.

## What could be shared

Three foundations are natural candidates:

- **A common way of describing a governed journey,** covering its outcome, accountable owner, consent, authority, human placements and recourse, so that good designs can travel between states and programmes.
- **Reusable trust capabilities,** such as durable consent, bounded delegation, checks before consequential actions, traceability and recourse routing, which programmes can inherit rather than rebuild.
- **Shared human capacity:** training, methods and communities of practice for the roles described in Section 7.

Each can start small, for example with two states reusing a reviewed journey design.

## A neutral steward for the shared layer

None of this requires a central platform. Ministries and states keep their journeys and decisions; what is shared is the trust layer, not control of the programme. Whoever maintains that layer gains quiet influence over how every journey behaves, so it is best held by a neutral, publicly accountable steward with no operational stake in any programme, while sectoral regulators keep their oversight. India's digital public infrastructure offers working precedents, where a steward maintains shared foundations while institutions keep their own customers, decisions and liabilities (RIS, 2024; MOSIP).

## Defined by purpose, not technology

The technologies of agent coordination are young and changing quickly. The shared layer is better defined by what it guarantees: what consent covers, who may approve, where risk changes, what must be traceable and how recourse begins. A layer built on purpose is more likely to outlast this season's technical fashion.

The note closes with the questions that seem most worth exploring together.

# 9 An Invitation

This note began with a parent, a caregiver, a pensioner and a farmer carrying the coordination that their services could not. It has described how that burden might move: to systems that hold the journey, within boundaries that hold the systems, under people who hold the boundaries. Models will keep improving on their own schedule. The institutional work needs deliberate effort.

Several questions now seem ready for collective attention, each larger than any one institution.

**Which journeys should we start with?** The best early candidates matter greatly to citizens, have a clear structure and can be judged by their outcomes: completion, repeated loops, time to resolution and silent stalls.

**Who should approve a governed journey, and how?** A review needs a clear owner, a defined scope and a way to share its findings across states and sectors.

**How will we build and support the people who keep journeys safe?** The roles in Section 7 exist in fragments today; they need training, tools and a career path.

**Which safeguards should be shared, and who should hold them?** These are choices for governments, domain institutions, civil society and technologists to make together.

**When should an agent be allowed to act on its own?** Today, every consequential step stays with a person. Any change should wait until the impact of each action is provably limited, every action remains traceable, and automatic checks are as dependable as the human decision they replace.

These questions are already drawing researchers, practitioners and national platforms together. The authors offer this note as a lens and a vocabulary to help frame them, expecting both to improve through use. Disagreement, refinement and evidence from real journeys will only make the proposition stronger.

India taught the world that identity, payments and data exchange could be public goods. This note asks whether the journey between a citizen and her outcome can be next. The industry has learned to make agents act; the harder and more valuable art is making them stay.

# Annex A: Six journeys used to test the argument across domains

The journeys below were used as a cross-domain test of the argument. They are not implementation specifications. Their value is the recurrence they reveal: long duration, exceptions, fragmented authority, changing risk and the need for structured human judgement.

| Journey | Companion role | Dominant exceptions | Human anchor |
|---|---|---|---|
| **RTE education inclusion** | Holds the parent journey from eligibility and document readiness through lottery, verification and school confirmation. | Name or date mismatch; post-lottery verification; school refusal; missed deadline. | Citizen Support Facilitator; district education officer or SCPCR (State Commission for Protection of Child Rights). |
| **Child immunisation continuum** | Carries the child schedule, reminders, catch-up planning and record continuity across the full cycle. | Migration; missed doses; session change; hesitancy; adverse-event signal. | ANM (auxiliary nurse midwife) or health worker; supervisory medical officer; programme support specialist. |
| **Maternal health** | Coordinates visits, tests, referral readiness and danger-sign escalation through pregnancy. | Risk reclassification; missed test; transport or referral failure; facility change. | ANM; clinician or medical officer; referral coordinator. |
| **ASHA work coordination** | Creates a coherent priority and follow-up view while respecting official-system authority. | Conflicting priorities; missing data; household vulnerability; no direct write authority. | ASHA; authorised medical officer; MIS support specialist. |
| **Dignity in pension** | Holds life verification, payment continuity and grievance resolution through an annual cycle. | Life-certificate failure; bank mismatch; silent payment stoppage; unresolved grievance. | Assisted-service facilitator; pension escalation officer; compliance reviewer. |
| **Farmer insurance claim** | Carries claim initiation, evidence readiness, verification, settlement and dispute. | Land or identity mismatch; contested assessment; missing evidence; settlement delay. | Facilitator; agriculture or revenue verifier; claims escalation coordinator. |

***Table 3**. Cross-domain recurrence is the evidence for a journey-level architecture and governance lens.*

Source note: the illustrations are synthetic composites based on analysis of public-service workflows and are intended to expose recurring design and governance patterns rather than describe a specific deployment.

# References


### AI governance and the agentic frontier

Ministry of Electronics and Information Technology. 2025. India AI Governance Guidelines: Enabling Safe and Trusted AI Innovation. Government of India. https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc2025115685601.pdf

Ministry of Electronics and Information Technology. 2025. IndiaAI Mission: Safe and Trusted AI Pillar. Government of India. https://indiaai.gov.in/mission

NITI Aayog. 2025. Roadmap on AI for Inclusive Societal Development. Government of India. https://niti.gov.in/sites/default/files/2025-10/Roadmap_On_AI_for_Inclusive_Societal_Development.pdf

UK Government. 2025. AI Insights: Agentic AI. https://www.gov.uk/government/publications/ai-insights/ai-insights-agentic-ai-html

Schmitz, C., Rystrøm, J. and Batzner, J. 2025. Oversight Structures for Agentic AI in Public-Sector Organizations. Proceedings of the 1st Workshop for Research on Agent Language Models (REALM 2025), ACL. https://arxiv.org/abs/2506.04836

### Digital public infrastructure

Research and Information System for Developing Countries. 2024. Foundations of Digital Public Infrastructure: A Handbook. New Delhi: RIS. https://ris.org.in/sites/default/files/Publication/DPIHandbook.pdf

MOSIP. Modular Open-Source Identity Platform. International Institute of Information Technology Bangalore. https://www.mosip.io

### Data protection, consent and accountability

Ministry of Electronics and Information Technology. 2023. The Digital Personal Data Protection Act, 2023. Government of India. https://www.indiacode.nic.in/handle/123456789/22037?view_type=browse

Ministry of Electronics and Information Technology. 2025. The Digital Personal Data Protection Rules, 2025. Gazette of India, notified 13 November 2025. Government of India. https://egazette.gov.in/WriteReadData/2025/267650.pdf

### Generative and agentic AI governance, traceability and observability

National Institute of Standards and Technology. 2023. Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. U.S. Department of Commerce. https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf

National Institute of Standards and Technology. 2024. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1. U.S. Department of Commerce. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

European Parliament and Council of the European Union. 2024. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), Article 12 (record-keeping) and Article 14 (human oversight). Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Wang, Y., et al. 2026. From Agent Traces to Trust: A Survey of Evidence Tracing and Execution Provenance in LLM Agents. arXiv preprint. https://arxiv.org/abs/2606.04990

### Human oversight and the journey workforce

McKinsey and Company. 2025. Superagency in the Workplace: Empowering People to Unlock AI's Full Potential. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work

Georgetown Center for Security and Emerging Technology. 2024. AI and the Future of Workforce Training. https://cset.georgetown.edu/publication/ai-and-the-future-of-workforce-training/

IndiaAI. 2026. Human Capital Working Group: Advancing Equitable Skilling for an AI-Enabled Future of Work. Government of India. https://impact.indiaai.gov.in/working-groups/human-capital

### The agentic journey and long-horizon autonomous agents

Anthropic. 2024. Building Effective Agents. https://www.anthropic.com/engineering/building-effective-agents

Anthropic. 2025a. Effective Context Engineering for AI Agents. https://www.anthropic.com/engineering/effective-context-engineering-for-ai-agents

Anthropic. 2025b. Effective Harnesses for Long-Running Agents. https://www.anthropic.com/engineering/effective-harnesses-for-long-running-agents

Du, P. 2026. Memory for Autonomous LLM Agents: Mechanisms, Evaluation, and Emerging Frontiers. arXiv preprint. https://arxiv.org/abs/2603.07670

Gemini Team, Google DeepMind. 2024. Gemini 1.5: Unlocking Multimodal Understanding across Millions of Tokens of Context. arXiv preprint. https://arxiv.org/abs/2403.05530

### Illustrative citizen journeys: underlying programmes

Ministry of Education, Department of School Education and Literacy. Right to Education Framework and Guidelines. Government of India. https://dsel.education.gov.in/rte

Ministry of Health and Family Welfare. 2023. Universal Immunisation Programme Guidelines and National Immunisation Schedule. Government of India. https://www.mohfw.gov.in/immunisation

Ministry of Health and Family Welfare. 2023. U-WIN Platform: Digitising the Universal Immunisation Programme. Government of India. https://uwin.mohfw.gov.in

Ministry of Health and Family Welfare. 2024. Reproductive and Child Health Portal: Maternal and Child Tracking. Government of India. https://rch.mohfw.gov.in

National Health Mission. ASHA Programme: Overview, Guidelines and Resources. Government of India. https://nhm.gov.in/index1.php?lang=1&level=1&sublinkid=150&lid=226

Ministry of Electronics and Information Technology. Jeevan Pramaan: Digital Life Certificate. Government of India. https://jeevanpramaan.gov.in

Ministry of Rural Development. National Social Assistance Programme. Government of India. https://nsap.nic.in

Ministry of Agriculture and Farmers' Welfare. Pradhan Mantri Fasal Bima Yojana. Government of India. https://pmfby.gov.in

Department of Land Resources. Digital India Land Records Modernization Programme. Government of India. https://dolr.gov.in/programmes-schemes/dilrmp-2/

# About IIIT-Bangalore and the Authors

IIIT-Bangalore bridges policy, development and technology through interdisciplinary teaching and applied research. The Centre for Information Technology and Public Policy (CITAPP) anchors the institute's human-centric approach to technology and governance. The Centre for Digital Public Infrastructure (CDPI) has led global work on population-scale identity and data systems, including MOSIP (Modular Open-Source Identity Platform), now adopted by over 25 countries. The Centre for Accessibility in the Global South (CAGS) advances research and practice on digital accessibility with national and international partners.

The A4I Lab (AI Innovation for Inclusion Initiative), a collaboration between IIIT-Bangalore and Microsoft, India, builds and deploys AI solutions for inclusion at population scale, including platforms in daily use in government school systems.

This note draws on analysis of citizen journeys across education, health, social protection and livelihoods. Together, these initiatives place the institute at the intersection where this note is written: digital design, inclusion and public policy at population scale.

## The Authors



**Ashish Srivastava** is the Chief Impact Officer at A4I Lab (AI Innovation for Inclusion Initiative), a collaboration between IIIT-Bangalore and Microsoft, India. He has over three decades of senior executive and technology leadership at renowned IT companies, spanning AI, Digital and Telecom. His recent work has centred on delivering population-scale digital solutions in the social sectors, including the agentic platforms whose journeys inform this note.

Ashish serves on the boards of leading institutions focused on technology, governance and societal impact, has served on government committees on health, data strategy and early child education, and has represented international delegations on AI collaborations. He received the IDEX AI Innovation Award (2019). He holds a Master's degree from IIIT-Bangalore, a B.Tech from IIT (BHU), and a post-graduate certification in Economics from the University of London.



**Prof. Amit Prakash** is Professor and Head of the Department of Digital Humanities and Societal Systems at IIIT-Bangalore, where he also serves as Convenor of CAGS. He established and leads CITAPP and the E-Health Research Centre at IIIT-B. His research explores the intersection of information systems and public policy, with a focus on public health, education and livelihood security; the emphasis on equity and inclusion in technology design is one reason the citizen, rather than the technology, sits at the centre of this note's framing. Prior to academia, he spent over 15 years in senior governance advisory roles at multinational consulting firms.

Prof. Prakash has chaired national and international committees on digital governance and is a resource person for the UN Asia-Pacific Centre for ICT and Development. In December 2023, he received the Marichi Award for transforming accessible education. Prof. Prakash holds a doctoral degree from IIM Bangalore and a B.Tech from IIT Roorkee.


iit-b
ज्ञानमुत्तमम्