

National Agentic Technology Infrastructure

Architecting the Agent-Native State — how the world's leading states are extending the digital-government canon to a workforce that reasons and acts.

Fewer than ~10 nations have wired the agentic foundation onto a healthy digital state — and perhaps two or three have done it well. That gap is the story of the agentic decade.

ABOUT THIS REPORT · HOW TO READ IT

The agentic state, evidenced

The European Nexus for Strategic Intelligence (ENSI) builds decision-grade intelligence on how nations architect technology. This issue maps the next layer of the state: the public infrastructure that lets **AI agents** — software that perceives, reasons and acts within bounded authority — do useful, safe, accountable work at population scale. It is not an AI-strategy PDF and not a chatbot procurement; it is a reading of the agentic literature against the national-architecture canon, resolved into a buildable stack.

The evidence base

This report is grounded in a verified collection of **142 downloaded PDFs and ~174 catalogued sources** spanning **12 folders** — from national architectures (Estonia, Singapore, India, the UAE) and the digital-state frameworks (eIDAS, NIST, FedRAMP) to the agentic corpus: the *Agentic State* vision paper, NASCIO 2026, OpenAI's governing-practices, the MCP and A2A protocols, and the agent-systems research (ReAct, RAG, Reflexion, MemGPT, τ -bench). Every number traces to a source; agentic adoption figures the literature could not count precisely are labelled **(est.)**.

Three reading depths

- ▶ **90 seconds** — this front matter: cover, summary, KPI dashboard, the one-diagram thesis and twelve findings.
- ▶ **15 minutes** — add the framework, the cross-country landscape and the ranked architectures.
- ▶ **Full read** — the 50 principles, 50 components and ten archetypes, with the cases, evidence and appendix.

The visual system — a legend

Magenta is the one accent: emphasis and the load-bearing layers. Green marks data-positive, coral marks an alert. Two coding systems run through every page — the five **priority tiers** and the eight **stack layers**.

PRIORITY TIERS — HOW ESSENTIAL A COMPONENT IS

- P0** Bedrock — no agentic state without it
- P1** Core — defines a serious agentic state
- P2** Differentiator — separates the leaders
- P3** Advanced — the maturing edge
- P4** Frontier — the experimental frontier

THE EIGHT STACK LAYERS — WHERE A COMPONENT LIVES

- L1** Compute & Models
- L2** Data, Knowledge & Memory
- L3** Tools & Integration
- L4** Trust & Control · cross-cutting
- L5** Agent Runtime
- L6** Orchestration
- L7** Experience & Mission
- L8** Governance Spine · cross-cutting

Source chips (e.g. **NASCIO 2026**) close findings and tables; **bold** renders magenta throughout and always marks the thing that matters on the page.

CONTENTS

What is inside

MOVEMENT A – THE 90-SECOND READ

01	Cover	1
02	About ENSI · how to read · the visual legend	2
03	Contents	3
04	Executive summary — the spine in six paragraphs	4
05	The dashboard — twelve headline indicators	6
06	The thesis in one diagram — the 8-layer stack	8
07	Twelve key findings	9

MOVEMENT B – FRAMEWORK & LANDSCAPE

08	The reframe — why "buy an AI" is the wrong question	11
09	From the digital state to the agentic state	13
10	The 8-layer National Agentic Technology Stack, explained	15
11	The global landscape & readiness	18
12	How we rank — priority tiers & the scoring lens	20

MOVEMENT C – PRINCIPLES, COMPONENTS & ARCHITECTURES

13	The 50 principles of the agent-native state	21
14	The 50 components — a bill of materials	30
15	The adoption curve — where nations fall off	37
16	The 10 best agentic architecture types, ranked	39
17	The optimal national agentic stack — the reference design	46

MOVEMENT D – EVIDENCE & CASES

18	Who's closest — the agentic-state scorecard	49
19	Seven nation case studies	50
20	Cross-cutting — sequence, failure modes & recommendations	57
21	Risks, counter-arguments & outlook to 2030	59

MOVEMENT E – BACK MATTER

22	Methodology & integrity notes	61
23	The 174-source evidence base — annotated catalogue	62
24	Back cover — the order is the strategy	64

Three depths: 90 seconds — pp. 1–10 · 15 minutes — add the framework, adoption curve & optimal stack · Full read — the 50 principles, 50 components, 10 archetypes & 7 cases. Every component carries a priority (**P0–P4**) and an adoption figure; agentic estimates are labelled (**est.**).

EXECUTIVE SUMMARY • 1 OF 2

The agentic state is the digital state, extended

For thirty years the organising metaphor of national technology was the *digital service* — a form turned into a web page, a registry turned into an API. The best states built that layer well: Estonia's X-Road, India's identity-payments-data stack, Singapore's government-as-a-platform, the EU's interoperability frameworks. That achievement is now the floor, not the ceiling. A new substrate has arrived — AI agents — and NASCIO frames the shift bluntly as the move from "AI that assists" to "AI that acts."

National Agentic Technology Infrastructure (NATI) is the public infrastructure a nation must build so that agents can act usefully, safely and accountably at population scale. It is a **stack** — compute, models, data, memory, tools, identity, orchestration, governance — re-conceived around a worker that reasons and acts. The central claim of this report is that the digital state and the agentic state are not rivals: **the agentic state is what you build on top of a working digital state, and it fails without one**. An agent is only as good as the identity it can verify, the registry it can trust, the tool it can call, and the human who can stop it.

The optimal design at this stage is clear and convergent. It is a **sovereign agent platform** provisioned centrally for all of government, on which every agent is a **grounded, reflective reason-act-verify loop** with governed memory, reaching the world through an **MCP tool layer** and other agents through an **A2A federation**, all running under **supervised autonomy** — bounded action, mandatory human gates, agent identity, authenticated delegation, full observability, a tested kill-switch — evaluated continuously against a mandatory eval harness and a binding national risk framework, and sitting atop a healthy digital state.

The win is not any single exotic pattern. It is having **all of the agentic P0-P1 components wired together onto a working digital foundation** — which fewer than ~10 nations have done, and perhaps two or three have done well. The decisive strategic fact is that the agentic race is being run *on the foundations laid by the digital race*: nations that built X-Road-class interoperability, population-scale identity and clean registries already hold half the agentic P0 stack, while those that skipped the digital foundation will find that no amount of agent procurement substitutes for it.

THE SPINE — THE CLAIM THIS WHOLE REPORT DEFENDS

The agentic state is the digital state's canon — identity, registries, federation, interoperability, governance, oversight — extended to a worker that reasons and acts. The win is having the agentic P0-P1 layer wired onto a healthy digital foundation, which fewer than ~10 nations have done.

The pages that follow translate this into a buildable architecture: a priority model from P0 Bedrock to P4 Frontier, an eight-layer stack, fifty principles, fifty components and ten archetypes — and a scorecard of who, today, is closest to the optimal design.

EXECUTIVE SUMMARY · 2 OF 2

What separates a thinking state from a demo

Two paragraphs of consequence, then the five things a decision-maker must internalise.

The defining risk of agents is **confident error**; the defining requirement of public administration is **defensible decisions**. Grounding resolves the tension — retrieval from authoritative registries, with citations, turns an agent's output into something it can prove, so ungrounded agents on official matters automate unaccountability. And because models are replaced every few months, a continuous **evaluation harness** measuring capability and consistency is the difference between a controlled rollout and a silent regression — τ -bench's finding that agents often fail on repeated trials is the single most important deployment caution in the evidence base.

The other half of the answer is governance as architecture. For software that *acts*, identity, authenticated delegation, least-privilege scopes, human approval gates and a tested kill-switch are not a brake bolted on afterwards — they are structural members of the design. Supervised autonomy is the only pattern present in *every* credible governance source, and the maturity discipline — climb, do not leap — is how a state earns autonomy as reliability is proven. The order is the strategy.

FIVE THINGS TO KNOW

- 1 It's continuity, not rupture.** The agentic state inherits the digital canon and extends it to a worker that acts — it does not replace it.
- 2 The gap is the story.** Fewer than ~10 nations have the agentic P0–P1 layer wired onto a healthy digital foundation; two or three have done it well.
- 3 Build a platform, not pilots.** One shared sovereign platform operationalises safety once and inherits it everywhere; a thousand pilots cannot be governed.
- 4 Earn the AI layer.** Intelligence is the capstone of identity → data → interoperability — resource the lower layers first or automate dysfunction.
- 5 Supervised autonomy is the architecture.** Bounded action, human gates, agent identity and a kill-switch make autonomous action legally survivable.

No nation has the full optimal stack. The frontier is a composite — the US's platform, Estonia's digital foundation and framework, the UK's assurance discipline, Singapore's governance toolkit, the EU's binding risk law, and the UAE's mission ambition. The report's task is to name the components, rank the architectures, and show the assembly order.

~10

nations with the agentic P0–P1 layer on a healthy digital state (est.)

8+

US states already running agentic tools in production — NASCIO 2026

50

design principles · 50 components · 10 architecture archetypes

THE DASHBOARD • 1 OF 2

The agentic state in twelve numbers

The foundation and the frontier, read together. The connectivity, identity and cyber numbers are hard global figures; the agentic-adoption tiles are labelled **(est.)** where the literature could not count precisely.

5.5^{bn}

people online in 2024
— 68% of humanity;
2.6bn remain offline

ITU FACTS & FIGURES
2024

96%

of the world covered
by a mobile-
broadband network

ITU 2024

84/4%

5G population
coverage — high-
income vs low-income
countries

ITU 2024

139/194

nations with a national
CIRT / CERT; 55 still
have none

ITU GCI 2024

~850^m

people still lack any
official ID — the
agent's root of trust

WORLD BANK ID4D 2021

198

economies scored on
GovTech maturity
across 48 indicators

WORLD BANK GTMI 2022

8⁺

US states already
running agentic tools
in production

NASCIO 2026

~10

nations with the full
agentic P0–P1 stack
on a healthy digital
state (est.)

ENSI SYNTHESIS — EST.

~5–15

nations standing up a
national agent
platform (est.)

ENSI SYNTHESIS — EST.

#1

network readiness —
USA; #2 Singapore,
#3 Finland

NRI 2024

~38

nations benchmarked
on digital-government
maturity

OECD DGI 2023

~20–25

nations with governed
model-access
programmes (est.)

ENSI SYNTHESIS — EST.

READING THE DASHBOARD

The top row is the **digital foundation** the agentic state is built on — broadband, coverage, 5G, cyber response. The lower rows are the **agentic frontier**, where hard counts thin out and the synthesis estimates begin: the population of nations with a wired P0–P1 stack is small, and the platform builders are fewer still. Estimates are derived from the leaders documented in folders 08–12 and are explicitly indicative.

THE DASHBOARD • 2 OF 2

Who leads — and the adoption cliff


Two hero charts and a share view. The leaderboard is an **indicative composite** against the optimal stack (0–100, est.); the drop-off shows how adoption collapses as you climb the priority tiers.

FIG 1 • INDICATIVE / EST.




Who is closest to the optimal agentic stack. Indicative composite score (0–100, **est.**) blending platform, grounding, supervised-autonomy, identity/delegation and digital-foundation strength. No nation has the full stack — the leaders are complementary strengths, not a single winner. Source: ENSI synthesis of folders 08–12.

FIG 2 • EST.



The adoption cliff. Many states hold the P0 digital bedrock; the count of nations falls steeply as the priority tier rises, with only ~2 reaching the frontier (**est.**). The gap between P0 and P1 is the agentic decade's central opportunity.

FIG 3 • EST.



The ~10-nation share. Roughly 10 of 193 UN member states have wired the agentic foundation onto a healthy digital state — about **5% (est.)**. The agentic state is, today, a minority achievement.

CAVEAT — READ THE ESTIMATES AS INDICATIVE

No global index yet counts "agentic-ready" nations. FIG 1–3 are ENSI composites built from the documented leaders (USAi, the *Agentic State* framework, NASCIO's production states, NAIS 2.0, the EU AI Act) and are **directional, not measured**.

THE MASTER DIAGRAM

The 8-layer National Agentic Technology Stack

Read top to bottom. Magenta bands are the load-bearing layers; each carries its name, one-line role, and its P0 component. Everything sits on the digital-state foundation — remove it and the stack falls.

FIG 4 · THE REFERENCE ARCHITECTURE



The optimal National Agentic Technology Stack (2026). The agentic state is the digital state's canon — identity, registries, federation, governance, oversight — extended to a worker that reasons and acts; the win is the agentic P0–P1 layer wired onto a healthy digital foundation, which fewer than ~10 nations have done. Source: ENSI synthesis document 3 of 3.

KEY FINDINGS · 1 OF 2

Twelve findings — the first six

The load-bearing conclusions of the evidence base, each closing with its strongest source. Findings 1–6 cover the strategy and the architecture of trust.

- 1 **It is continuity, not rupture.** The agentic state inherits the digital state's canon — identity, registries, federation, interoperability, governance — and extends it to a worker that reasons and acts; it does not replace it. — Algorithmic State Architecture, arXiv 2503.08725
- 2 **Fewer than ~10 nations have closed the gap.** Only a handful have wired the agentic P0–P1 layer onto a healthy digital state, and perhaps two or three have done it well — the central strategic fact of the agentic decade. — ENSI synthesis (est.)
- 3 **Build a platform, not pilots.** The leaders provision a shared sovereign agent platform — model access, guardrails, observability and a tool catalogue — so every ministry inherits safety instead of re-failing at it. — GSA USAi · EU GenAI4EU
- 4 **Grounding is the antidote to confident error.** Every consequential agent must retrieve from authoritative sources and cite them; ungrounded generation on official matters is a defect, not a feature. — RAG, arXiv 2005.11401 · OECD/GPAI
- 5 **Supervised autonomy *is* the architecture.** Bounded action spaces, mandatory human gates, escalation paths and a tested kill-switch are the only pattern present in every credible governance source — a structural member, not a brake. — EU AI Act · NASCIO 2026 · UK AI Playbook
- 6 **Every agent needs a verifiable identity and scoped delegation.** An action you cannot attribute is an action you cannot govern; agents act only under least-privilege, time-bounded, auditable, revocable delegation. — Visibility into AI Agents, arXiv 2401.13138 · MIT, arXiv 2501.09674

THE THROUGHLINE

Findings 1–6 establish the strategy (continuity, the gap, the platform) and the architecture of trust (grounding, supervised autonomy, agent identity). Findings 7–12, overleaf, cover how agents reach the world, how they are proven reliable, and where the adoption curve breaks.

KEY FINDINGS · 2 OF 2

Twelve findings — the second six

Findings 7–12 cover integration, reliability, data readiness, the leaders, the cliff, and the rule that governs the whole sequence.

- 7 **MCP + A2A are federation reborn at the agent layer.** The Model Context Protocol (agent-to-tool) and Agent2Agent (agent-to-agent) reproduce X-Road's federation for agents — integrate through open standards, preserve agency ownership, avoid a central black box.
— Anthropic MCP · A2A → Linux Foundation
- 8 **No deployment without an eval harness — and continuous re-evaluation.** Agents must be measured for capability *and* consistency; τ -bench's pass^k shows they often fail on repeated trials, and every model change demands re-qualification. — τ -bench · AgentBench · GAIA
- 9 **Data readiness is a precondition, not a parallel workstream.** An agent automates whatever quality it is fed, at scale; agent-ready, clean, machine-readable registries are the new precondition the way machine-readable data once was. — OECD/GPAI government-data report
- 10 **No nation has the full stack — the frontier is a composite.** The US's platform, Estonia's foundation and framework, the UK's assurance, Singapore's governance toolkit, the EU's risk law and the UAE's mission ambition. — ENSI scorecard, folders 08–12
- 11 **The adoption curve falls off a cliff above P0.** Many states hold the P0 digital bedrock; the count collapses through P1–P4, with only ~2 reaching the frontier — the gap between P0 and P1 is the opportunity.
— ENSI components analysis (est.)
- 12 **You earn the AI layer; you do not skip to it.** Intelligence is the capstone of Digital-Public-Infrastructure → Data → Interoperability; resource the lower layers first or you automate dysfunction at machine speed.
— Algorithmic State Architecture, arXiv 2503.08725

FROM FINDINGS TO THE BUILD

These twelve findings are a dependency graph, not a menu. The remainder of the report walks them in order — the framework and landscape (Movement B), the principles, components and architectures (Movement C), the cases and recommendations (Movement D), and the risks and evidence base (Movement E) — so a nation can see not just *what* the agentic state is, but the sequence in which it is actually built.

“

The agentic race is being run on the foundations laid by the digital race — no amount of agent procurement substitutes for the identity, registries and interoperability beneath it.

— ENSI, NATIONAL AGENTIC TECHNOLOGY INFRASTRUCTURE · SYNTHESIS 3 OF 3

MOVEMENT B · THE REFRAME

01 Why "buy an AI" is the wrong question

Almost every government that has put a budget line against "agentic AI" has framed it as a purchase: a chatbot here, a copilot there. That framing is the costliest mistake a nation can make — because an agent that reasons and *acts* on citizens is not a product you procure. It is infrastructure you provision.

For thirty years the organising metaphor of national technology was the **digital service** — a form turned into a web page, a registry turned into an API. The contrarian core of this report is that the next layer breaks that procurement logic. NASCIO's 2026 report frames the shift as the move from "*AI that assists*" to "*AI that acts*." Software that "perceives, reasons, and takes autonomous action within defined boundaries" is not a feature bolted onto a service. It is a new **factor**

of production — a class of digital labour that converts policy into daily service.

And labour is provisioned, not bought off a shelf. The US GSA's **USAI** is a government-wide platform — chat, API and an evaluation console over multiple models — offered to every agency at no marginal cost; the EU's **GenAI4EU** plays the same role. Neither is "a chatbot the government bought." Each is a shared substrate on which safety, identity, evaluation and cost-control are operationalised *once* and inherited everywhere.

The category error, stated plainly

A chatbot answers; an agent acts. The moment software can move money, change a legal status, or take an irreversible step, it inherits the entire governance apparatus of the state — identity, authorisation, audit, liability, oversight.

You cannot procure that as a SKU. The question is therefore never "*which AI do we buy?*" It is "*what infrastructure must we provision so that agents can do useful, safe, accountable work at population scale?*"

“

The agentic state is the digital state's canon — identity, registries, federation, interoperability, governance, oversight — extended to a worker that reasons and acts.

— **THE SPINE**, NATIONAL AGENTIC TECHNOLOGY INFRASTRUCTURE

8+

US states run agentic tools in production (NASCIO 2026)

1

government-wide platform — USAi — over many models, free to agencies

~10

nations with the agentic P0 layer wired onto a digital foundation

200

real government AI use-cases surveyed by the OECD

MOVEMENT B · THE REFRAME · CONTINUED

Myth versus mechanism

Every dollar mis-spent on agentic AI can be traced to one of a handful of category errors. Each myth is intuitive, procurement-friendly and wrong. Each mechanism is the infrastructure pattern the leaders actually adopted instead.

THE MYTH (THE WRONG QUESTION)	THE MECHANISM (WHAT THE LEADERS DO INSTEAD)
Buy a chatbot for each agency.	Provision one sovereign agent platform — shared model access, shared guardrails, shared observability, shared tool catalogue — that every ministry builds on (USAI, GenAI4EU). Safety and capability are inherited, not re-purchased.
AI that assists.	AI that acts, under supervised autonomy — bounded action spaces, mandatory human-approval gates, a tested kill-switch. Autonomy is earned level-by-level up a maturity model, not switched on (NASCIO five-phase model).
More models = more capability.	Tools + grounded data + identity = capability. An agent without tools is a text generator; an agent with the right tools, retrieving from authoritative registries, acting under a verifiable identity, is a worker (Toolformer; RAG; MCP).
Agents replace the digital state.	Agents are the digital state's canon, extended. They amplify whatever they stand on — including bad data and broken process. They fail without the eID, registries, data-exchange and security beneath them.
It's an AI strategy (a PDF).	It's an operating stack — compute, models, data, memory, tools, identity, orchestration, governance — provisioned, instrumented and measured by resolved outcomes, the way the EU Digital Decade reports binding KPIs.

The vanity metric

The number of agents *deployed* measures nothing. Count what matters: a resolved citizen outcome, a civil-servant hour returned, a cost-per-resolution that falls. Instrument the platform for task-completion, reliability under repetition, and time saved — and report it like a binding target, not a press release.

The fastest lever

Procurement, not pilots, moves adoption. The US OneGov strategy buys for "government as one customer," with agentic deals struck centrally and FedRAMP 20x fast-authorisation clearing the path. Buy agent *capabilities and tokens* with pre-cleared models — not bespoke per-agency contracts that re-litigate safety each time.

THE ONE-LINE TEST

If the procurement document names a product, it is the wrong question. If it names a *capability provisioned once and governed for all of government*, it is the right one. Agents are infrastructure, not a feature — and infrastructure is provisioned, instrumented and governed, never simply bought.

MOVEMENT B • CONTINUITY

02 From digital state to agentic state

The agentic state and the digital state are not rivals. The agentic state is what you build *on top of* a working digital state — and it fails without one. An agent is only as good as the identity it can verify, the registry it can trust, the tool it can call, and the human who can stop it.

The *Algorithmic State Architecture* (arXiv 2503.08725) models the state as four stacked layers — **Digital Public Infrastructure → Data → Interoperability → Intelligence** — with the intelligence (agentic) layer explicitly sitting *on top of* the other three. You do not skip to it; you earn it. The strategic consequence is inconvenient: nations that built X-Road-class interoperability, population-scale identity and clean registries start the agentic race with half the P0 stack already in place. Those that skipped the digital foundation will

find no amount of agent procurement substitutes for it.

This is why the FDA's agency-wide agentic premarket-review rollout depends on *structured submission data*; why Estonia's *Agentic State* paper lists existing X-Road interoperability and a "sovereign tech stack" as enablers. The continuity is mechanical, not rhetorical — each load-bearing principle of the digital canon has a direct agentic successor that inherits the parent's discipline rather than replacing it. Read bottom to top, the four layers are a dependency graph, not a menu.

L4 • INTELLIGENCE (AGENTIC)

Agents that reason and act — grounded, supervised, governed. The capstone, reached only after the three layers beneath it.

L3 • INTEROPERABILITY

X-Road-style signed data exchange, base registries as single source of truth. Federation reborn as MCP + A2A.

L2 • DATA

Clean, current, machine-readable, governed records. The ground truth an agent grounds against via RAG.

L1 • DIGITAL PUBLIC INFRASTRUCTURE

eID, payments rails, connectivity, sovereign cloud, zero-trust. The trust root every transaction depends on.

THE CONTINUITY THESIS

The agentic state is not a leap past the digital state. It is the digital state's canon — identity, registries, federation, interoperability, governance, oversight — extended to a worker that reasons and acts. The order is the strategy: L1–L3 first, intelligence last.

MOVEMENT B · CONTINUITY · THE INHERITANCE MAP

What each digital principle becomes

The agentic canon is the digital canon, line by line, with one world-changing clause added: the actor is now a worker that reasons and acts. Where a NATI principle has a direct digital ancestor, the synthesis flags it [↔ Pn]. The map below reads the agentic layer *against* the national-architecture canon.

DIGITAL-STATE PRINCIPLE / COMPONENT	ITS AGENTIC SUCCESSOR
eID + assurance levels (trusted identity at the root of every transaction)	Agent identity + authenticated delegation — every agent has a verifiable identity; a citizen or official grants it scoped, time-bounded, revocable authority (OAuth/OIDC extended, arXiv 2501.09674).
Base registries (single source of truth)	Agent-ready registries / RAG ground truth — clean, queryable systems of record the agent grounds against; what it generates is a derived view, never an alternative truth.
X-Road federation (autonomous nodes, signed point-to-point exchange)	MCP + A2A — agent-to-tool (Model Context Protocol) and agent-to-agent (Agent2Agent) standards: federation reborn at the agent layer.
Interoperability framework (EIF; reusable building blocks; RIHA catalogue)	Tool & agent registries — a governed catalogue of approved tools and agents with capability metadata and assurance status (AGNTCY Agent Directory).
Zero-trust security (PDP/PEP, continuously verified, least privilege)	Action-scoping + kill-switch — per-agent least-privilege action sets, default-deny on irreversible operations, a tested emergency stop and rollback.
Once-only principle (don't re-ask for what the state already holds)	Governed tiered memory — working, long-term and episodic case memory, governed as a regulated personal-data store (purpose limits, retention, deletion).
Central digital agency (binding reference architecture, law, targets)	Central agent authority — one accountable owner of the national agent platform, under law defining permissible autonomous action and liability.
Tamper-evident logging (Estonia's KSI ledger)	Agent observability + immutable logs — queryable, tamper-evident records of reasoning, tool calls and actions: the precondition for any audit.

8

digital principles with a direct agentic successor

NATI SYNTHESIS

4

layers in the Algorithmic State Architecture

arXiv 2503.08725

1

added clause: a worker that reasons and acts

THE SPINE

“
The optimal national agentic architecture is the same sentence as the digital one — a federated, API-first, building-block platform with a trusted identity at its root and a signed exchange layer instead of a central database — now staffed by governed, grounded, supervised agents that reason and act on top of it.

— **SYNTHESIS 3**, THE OPTIMAL NATIONAL AGENTIC STACK

MOVEMENT B · THE FRAMEWORK

03 The 8-layer National Agentic Stack

The reference design is one stack of eight layers, sitting atop a healthy digital state. Two layers — Trust & Control (L4) and the Governance Spine (L8) — are cross-cutting. The win is not any single exotic pattern; it is having all eight wired together.

FIG B-1

LAYER	WHAT IT IS	TIER
L8	Governance Spine (cross-cutting) — binding risk framework · liability law · central agent authority · eval harness · oversight	P0
L7	Experience & Mission — citizen super-assistant · civil-servant copilots · policy twin · proactive life-event services	P0
L6	Orchestration — durable stateful workflows · orchestrator-worker · A2A federation · human-approval nodes	P1
L5	Agent Runtime — reason-act-verify loop · reflection · planning · skill library	P0
L4	Trust & Control (cross-cutting) — agent identity · delegation · action scopes · kill-switch · observability · gateways	P0
L3	Tools & Integration — MCP servers over every system · tool/agent registry · computer-use legacy bridge	P0
L2	Data, Knowledge & Memory — agent-ready registries · RAG · vector+graph · context engineering · governed memory	P0
L1	Compute & Models — governed frontier + sovereign models · model gateway · managed inference · cost routing	P0

The 8-layer stack. Read bottom-up. ▼ Sits atop the digital state: broadband, eID, base registries, X-Road-style exchange, sovereign cloud, zero-trust — the agentic stack fails without it.

L1–L3 — the engine room, the knowledge, the reach

L1 · COMPUTE & MODELS

Role. The engine every agent runs on. **P0 component:** governed frontier-model access (#1). **Signal:** ~15–25 nations run a governed national model-access programme; USAi abstracts many models behind one gateway. **P0 bedrock**

L2 · DATA, KNOWLEDGE & MEMORY

Role. What agents know and remember. **P0 component:** authoritative, agent-ready registries (#8) + RAG layer (#10). **Signal:** interoperable base registries in ~40 nations; *agent-ready* in fewer. **P0 bedrock**

L3 · TOOLS & INTEGRATION

Role. How agents reach the world. **P0 component:** MCP tool layer (#41) — the "USB-C for AI." **Signal:** standard ~18 months old; early adoption in a handful of leaders. **P0 bedrock**

MOVEMENT B · THE FRAMEWORK · CONTINUED

L4–L6 — control, runtime, orchestration

The middle of the stack is where autonomy is made survivable. L4 (Trust & Control) is cross-cutting: nothing autonomous ships before it. L5 builds the worker; L6 makes it run for hours and across agencies.

L4 · Trust & Control P0 bedrock · cross-cutting

Role. The structural member that makes autonomous action legally survivable. **P0 component:** agent identity (#33) + authenticated delegation (#34) + secure gateways (#40), with action-space constraints, a tested kill-switch (#36) and immutable observability (#37). **Signal:** production-grade national agent-identity exists in *single digits of nations* (est.). *This is the gate the whole stack turns on.*

L5 · Agent Runtime P0 bedrock

Role. How the worker itself works — the atomic unit of agency. **P0 component:** the standard *reason-act-verify* loop (#23), from ReAct and the Claude Agent SDK, with reflection (#24, Reflexion) and planning (#25) layered on. The verify step separates a reliable agent from a confident-but-wrong one. **Signal:** the de-facto standard in every serious build (est.).

L6 · Orchestration P1 core

Role. Turning a fragile prompt-chain into a durable process that survives crashes and hand-offs. **P1 component:** a stateful orchestration runtime (#27, LangGraph/AutoGen) with orchestrator-worker decomposition (#28), A2A inter-agency federation (#42) and human-approval nodes. **Signal:** standard tooling embedded in leading platforms; multi-agent orchestration a front-runner-only differentiator (est.).

THE LOAD-BEARING WALL

Supervised autonomy — bounded action, mandatory human gates, agent identity, authenticated delegation, full observability, a tested kill-switch — is the only pattern that appears in *every* credible governance source: the UK AI Playbook, NASCIO, the EU AI Act's high-risk requirements, OpenAI's practices. For a state it is not a constraint on the architecture. It *is* the architecture.

MOVEMENT B · THE FRAMEWORK · CONTINUED

L7–L8 — mission and the governance spine

The top of the stack is where citizens meet the agentic state, and the spine that runs through all of it. L7 is built last and continuously; L8 in parallel from day one. An L7 mission deployed without L1–L6 beneath it is the most dangerous configuration in the design.

L7 · Experience & Mission P0 bedrock


Role. The visible agentic state — where policy becomes daily service. **P0 component:** citizen super-assistants, civil-servant copilots, a national policy twin and proactive life-event services — all *bounded and human-gated*. **Signal:** the UAE's agents-in-service vision; Singapore's NAIS 2.0; the Tony Blair Institute's "Reimagined State." **Where to start:** the Alan Turing Institute found ~41% of GB public-sector working time is supportable by generative AI (49% in education) — target the highest-volume, highest-burden, *lowest-risk* functions first.

L8 · Governance Spine P0 bedrock · cross-cutting

Role. Governance is architecture; for agents that *act*, it is the load-bearing wall. **P0 components:** a binding national risk framework (#48 — NIST AI RMF / EU AI Act / Singapore's Model Gov Framework) and a mandatory eval harness (#47 — AgentBench, GAIA, τ -bench's pass^k metric), plus a liability chain (#38) and an oversight institution (#49). **Signal:** a *binding* framework applied to agents exists in ~10–20 nations; mandatory eval regimes in very few (est.). τ -bench's finding that agents fail on repeated-trial consistency is the single most important deployment caution in the collection.

The component chips, by tier

The 50 components map onto the eight layers. The distinctively *agentic* P0 set is the gate almost every nation fails:



THE BUILD SEQUENCE, IN ORDER

Pre-requisite: fix the digital state. Then L1–L2 foundation → L3–L4 trust spine (*nothing autonomous ships before L4*) → L5–L6 runtime → L8 governance in parallel from day one → L7 mission last and continuously. The order is the strategy.

MOVEMENT B · THE GLOBAL LANDSCAPE

04 Where nations actually stand

Readiness for the agentic state is not one number; it is two. A nation's position is set by the maturity of its **digital foundation** (eID, registries, exchange, cloud) and its **agentic capacity** (platform, identity/delegation, eval, MCP). Plot the two and four quadrants appear — and almost everyone is in the wrong one.

FIG B-2

	LOW AGENTIC CAPACITY	HIGH AGENTIC CAPACITY
Strong digital foundation	Primed. Half the agentic P0 stack already built (registries, eID, exchange) but no platform, agent identity or eval yet. <i>Estonia, Denmark/Nordics, Korea, India.</i>	Agentic frontier. Foundation + the agentic P0–P1 layer wired on. Fewer than ~10 nations; perhaps 2–3 done well. <i>United States (federal), Singapore, UK, leading US states.</i>
Weak digital foundation	Behind the gate. Neither layer. Buying agents here yields demos, not infrastructure — you automate dysfunction at machine speed. <i>The long tail of the ~190 economies in the GTMI.</i>	Built on sand. Mission ambition and procurement energy without clean registries or identity beneath — the most dangerous quadrant. <i>The cautionary case for any nation that skips L1–L4.</i>

The 2×2 readiness typology. Axes: digital-foundation maturity (vertical) × agentic capacity (horizontal). The frontier quadrant is reached only by building the foundation first — the "built on sand" quadrant is reached by skipping it.

The empirical spine beneath the map is the global connectivity and government-tech census — the floor every agentic deployment stands on:

5.5bn

people online (68% of humanity); 2.6bn still offline — ITU 2024

96%

of the world covered by a mobile-broadband network — ITU 2024

84% / 4%

5G population coverage: high-income vs low-income — ITU 2024

139/194

nations have a national CIRT/CERT — ITU GCI 2024

~850m

people still lack any official ID — World Bank ID4D

198

economies scored on GovTech maturity
WORLD BANK GTMI

~38

nations benchmarked on digital-government maturity
OECD DGI 2023

~40

nations with interoperable digital base registries
FOLDER 04 (est.)

<10

nations with the agentic P0 layer on a digital foundation
NATI SYNTHESIS (est.)

MOVEMENT B · THE GLOBAL LANDSCAPE · READINESS

The readiness leaders — and what each is strongest at

Network readiness is the best available proxy for the digital foundation the agentic race is run on. The Network Readiness Index 2024 ranks 133 countries; its leaders are the same nations closest to the optimal agentic stack.

FIG B-3



Network Readiness Index 2024 (relative standing). Bar length = relative NRI rank among 133 countries; #1 USA, #2 Singapore, #3 Finland. The digital-readiness leaders are the pool from which the agentic frontier is drawn.

But readiness only describes the foundation. The agentic scorecard ranks each leader by the archetype it is strongest at — and names the gap that keeps even the front-runners short of the full optimal stack:

NATION / BLOC	STRONGEST ARCHETYPE	MAIN GAP
US (federal)	Sovereign platform	Coherent governance across a fast rollout (USAi + OneGov + FDA agents).
Estonia	Federation + capstone	Scaling the <i>Agentic State</i> framework to a production platform.
United Kingdom	Supervised autonomy	A unifying agent platform at USAi scale (AI Playbook + Turing).
Singapore	Platform + governance	Moving from copilots to bounded autonomy (NAIS 2.0 + AI Verify).
UAE	Mission / experience	Open governance + eval depth (Minister of AI; agents in visas).
EU (bloc)	Federation + governance	Speed — turning frameworks into deployed platforms (Apply AI).
Leading US states	Supervised autonomy	Shared platform + identity depth (8+ in production — NASCIO).

THE DECISIVE FACT

No nation has the full optimal stack. The frontier is a composite — the US's platform + Estonia's foundation + the UK's assurance discipline + Singapore's governance toolkit + the EU's binding risk law + the UAE's mission ambition. The agentic race is run on the foundations laid by the digital race.

MOVEMENT B · THE SCORING LENS

05 How we rank: priority tiers

Every component and architecture in this report is placed on a five-step priority ladder. The tiers are not a wish-list; they are a dependency graph. P0 is what you build first because without it there is no *safe* agentic action at all.

P0 Bedrock

P1 Core

P2 Differentiator

P3 Advanced

P4 Frontier

TIER	DEFINITION	WHO HAS IT
P0 · Bedrock	Without it there is no safe agentic action. Build first. <i>19 components.</i>	Inherited digital P0: dozens. Agentic P0: single digits to ~15 nations.
P1 · Core	Defines a serious agentic state; present in the leaders. <i>16 components.</i>	~5–20 — US, EU bloc, UK, Singapore, UAE, Estonia, leading US states.
P2 · Differentiator	Separates front-runners from followers. <i>9 components.</i>	~3–10 — GraphRAG, multi-agent, skill libraries, agent registries.
P3 · Advanced	Top-tier only. <i>4 components.</i>	~2–8 — edge inference, computer-use at scale, episodic memory.
P4 · Frontier	The fully agent-native state — intelligence as capstone.	~1–3 — the state that built every layer beneath it first.

The adoption-evidence method

Agentic infrastructure is two-to-three years old as a *governed national* concern, so there is no ITU-grade census of it. Every adoption figure in this report is therefore one of three kinds, always labelled — and the reader should weigh them accordingly.

1 · HARD SIGNALS

Drawn directly from sources: *8+ US states in production* (NASCIO 2026); the FDA's rollout; USAi & GenAI4EU exist; ITU/World-Bank census numbers. Cited to source, treated as fact.

2 · INHERITED DIGITAL COUNTS

For components the agentic state reuses from the digital state — identity, registries, cloud — counted via folder 04's indices (~40 nations with interoperable registries). Solid, but a *floor*: digital ≠ agent-ready.

3 · LABELLED ESTIMATES

Reasoned, order-of-magnitude bands marked (est.) — triangulated from the strategy literature. Never precise counts.

THE HONEST HEADLINE

On the foundational *digital* components agents depend on, dozens of nations qualify. On the distinctively *agentic* components — agent identity, delegation, eval harnesses, governed memory, the MCP layer, the national platform — the number with a governed, production-grade capability is in the **single digits to low tens**. That is precisely why the window is open.

The 50 Principles of the Agent-Native State

The convergent design canon that emerges from reading the agentic literature against the national-architecture literature — fifty rules organised into **seven families** (A Systemic & strategic · B Agent design & reasoning · C Models & compute · D Data & state · E Identity, trust & accountability · F Tools & integration · G Evaluation, safety & resilience) and tiered from  **Foundational** through  **Core** and  **Enabling** to  **Frontier**. They are not a menu but a dependency graph: earn the right to deploy, become a serious agentic state, separate from the pack, then reach the frontier.

How the whole agentic state is shaped

The agentic state is not an AI strategy PDF or a chatbot procurement; it is a governed expansion of the state's capacity to act, built deliberately on top of a working digital foundation and resourced like headcount.

- 1 **Treat agents as a new factor of production, not a feature.** ■ Stop bolting AI onto services; treat agents as *digital workers* that convert policy into daily service. National implication: plan agent capacity the way you plan headcount and budget, with workforce-transition planning attached.
- 2 **Build the agentic layer on top of a working digital state, never instead of one.** ■ Every credible deployment sits on healthy DPI — identity, registries, data exchange, APIs; agents amplify whatever they stand on, including dysfunction. National implication: get the P0–P2 digital components healthy *before* scaling autonomy, or you automate dysfunction at machine speed.
- 3 **Agents are infrastructure, not apps — provision them centrally.** ■ Leaders stand up shared agent *platforms* (GSA's USAi, the EU's GenAI4EU), not app-by-app buys. National implication: run one national platform so every ministry inherits shared models, guardrails, observability and a tool catalogue.
- 4 **Adopt a maturity model and sequence autonomy deliberately.** ■ NASCIO's five phases run assistive → context-aware → task automation → stateful workflows → adaptive agents; states climbed, they did not leap. National implication: publish a maturity model and gate each rise in autonomy on demonstrated reliability at the prior level.
- 5 **Start where the time is, not where the hype is.** ■ The Turing Institute found ~41% of GB public-sector working time is GenAI-supportable (49% in education). National implication: run a time-and-task analysis and target the highest-volume, highest-burden, lowest-risk functions first.
- 6 **Govern with binding law and a central authority.** ■ Architecture sticks only when backed by legislation and a central agency — doubly so for agents that take legally consequential *actions*. National implication: name one accountable authority operating under law that defines permissible autonomous action and liability.
- 7 **Design for measurable outcomes, not deployed models.** ■ Agents-deployed is a vanity metric; the unit of value is a resolved citizen outcome or a civil-servant hour returned. National implication: instrument the platform for task-completion, reliability, cost-per-resolution and time-saved, reported like binding KPIs.
- 8 **Sovereignty applies to agents as it does to data.** ■ An agent reasoning over national data and acting on citizens must run under national control. National implication: hold model, orchestration and memory layers to residency and operational-autonomy controls; keep a sovereign / open-weight option for the most sensitive work.
- 9 **Procurement is the fastest lever you have.** ■ OneGov "government as one customer" and FedRAMP 20x move adoption faster than any pilot. National implication: buy agent capabilities and tokens centrally, with pre-cleared models and pre-authorised tool integrations, not bespoke per-agency contracts.
- 10 **Plan for the AI layer to sit atop data and interoperability — earn it.** ■ The *Algorithmic State* models four layers (DPI · Data · Interoperability · Intelligence) with intelligence explicitly on top. National implication: treat the agentic layer as the capstone and resource the lower layers first — an agent with no clean data is a demo, not infrastructure.

THE LOAD-BEARING SEQUENCE

Family A contains the strategy's spine: P2 and P6 are non-negotiable ■ foundations — a working digital state and a governing authority. Skip them and everything above is an expensive demo.

10

principles · Family A

3

■ foundational

~41%

GB time GenAI-able

How the workers themselves work

Autonomy is a cost and a risk, spent only where dynamism is genuinely required. The atomic unit of agency is an inspectable reason–act–observe loop — and legibility is a governance asset, not a luxury.

- 11 **Prefer the simplest pattern that works; not everything needs an agent.** ■ Anthropic separates *workflows* (predefined code paths) from *agents* (LLMs directing their own process) and argues for the simplest composition. National implication: make "don't use an autonomous agent where a deterministic workflow suffices" the engineering default.
- 12 **The reason–act–observe loop is the atomic unit of agency.** ■ ReAct interleaves verbal reasoning, tool actions and environment observations so reasoning guides action and observation grounds reasoning. National implication: standardise reference designs on an explicit, inspectable loop rather than opaque end-to-end black boxes.
- 13 **Build the loop as gather–context → act → verify.** ■ The Claude Agent SDK's cycle adds a *verify* step that separates a reliable agent from a confident-but-wrong one. National implication: require a verification stage (self-check, test, or second-agent review) before any consequential or irreversible action is committed.
- 14 **Give agents reflection and self-correction.** ■ Reflexion's verbal self-critique stored in episodic memory improves later attempts without retraining. National implication: prefer architectures that detect and recover from their own errors — a system that catches its own mistake beats one that never doubts itself.
- 15 **Planning and decomposition beat single-shot generation on hard tasks.** ■ Chain-of-Thought unlocked multi-step reasoning; Tree-of-Thoughts adds search with lookahead and backtracking. National implication: reference agents for benefits, procurement and casework should plan and decompose explicitly, with the plan exposed for audit.

“

A loop you cannot read is a loop you cannot govern. Interpretability is not a nicety in the public sector — it is the precondition for accountability.

— **SYNTHESIS**, NATI PRINCIPLES (FAMILY B)

IN SHORT

The first half of Family B fixes the *shape* of a public-sector agent: simplest-pattern-first, an explicit ReAct loop, a mandatory verify step, self-correction, and visible planning. Two of these five (P11, P12) are ■ foundational — they govern whether an agent can be inspected and trusted at all.

5 / 10

Family B shown here

2

■ foundational

P11–P15

design fundamentals

Tools, memory and legibility

An agent without tools is a text generator; an agent with the right tools is a worker. Memory turns a chatbot into a colleague — and a colleague whose reasoning, plans and tool calls are all on the record.

- 16 Tool use is a first-class capability, learned not bolted on.**  Toolformer showed models learn *when* and *how* to call APIs natively; Gorilla and ToolLLM scaled it to thousands of real APIs. *National implication:* a platform's value is dominated by the breadth and quality of tools it can safely call — invest in the tool layer at least as much as in models.
- 17 Memory turns a chatbot into a colleague.**  Generative Agents' recency/importance/relevance memory stream and MemGPT's OS-style paging let an agent carry a case across sessions. *National implication:* design memory deliberately and govern it — an agent's long-term memory of citizens is itself a regulated data store.
- 18 Accumulate reusable skills, don't re-derive them.**  Voyager built an ever-growing library of self-verified, executable skills that compound — the agentic analogue of Korea's eGovFrame component libraries. *National implication:* catalogue, version and reuse validated agent skills across agencies rather than rebuilding the same casework agent.
- 19 Match the architecture to the task's irreversibility and ambiguity.**  The surveys show no single best design — the right autonomy, memory and planning depend on the task. *National implication:* publish a decision guide mapping task properties (reversibility, ambiguity, volume, risk) to recommended patterns, so teams choose by analysis, not fashion.
- 20 Keep the agent legible: expose reasoning, plans and tool calls.**  OpenAI lists legibility of agent reasoning as a baseline safety practice. *National implication:* mandate that production agents emit structured, retainable traces of reasoning, plans, tool calls and observations — legibility is both a debugging tool and an accountability record.

WHY THE TOOL LAYER WINS

P16 is the hinge of the whole stack: capability is dominated not by the model but by the catalogue of governed tools an agent can reach. It is why Family F (tools & integration) carries five   principles of its own — and why a national platform should out-invest its models with its connectors.

10

principles · Family B total

P16–P20

tools · memory · skills · legibility

1000s

APIs scaled by ToolLLM

The engine room

Agents are only as capable as the models beneath them, and frontier capability is concentrated in a handful of providers. Model access is now the agentic equivalent of energy security — to be secured, diversified and managed, never assumed.

- 21 **Treat model access as a strategic national input.** ■ Frontier capability is concentrated; the US AI Action Plan makes infrastructure a top pillar. *National implication:* secure governed, resilient access to frontier models (commercial *and* open-weight) with continuity plans for provider or jurisdictional disruption.
- 22 **Maintain model optionality and avoid single-model lock-in.** ■ USAi is a unified interface over multiple leading models, not a bet on one; capability and price shift monthly. *National implication:* abstract the model behind a stable interface so agencies route by task, cost and assurance, and the state can switch providers without re-engineering every agent.
- 23 **Keep a sovereign / open-weight option for the most sensitive work.** ■ Where data cannot leave national control, open-weight models on sovereign infrastructure are the fallback. *National implication:* sustain at least one in-country, controllable inference path — even if less capable — so the state is never wholly dependent on an external API for consequential action.
- 24 **Provision inference capacity as managed national infrastructure.** ■ Agentic workloads are inference-heavy and bursty — one task fans out into dozens of model calls. *National implication:* plan inference capacity, latency and cost as a managed shared service with quotas and prioritisation, not an unmanaged per-agency cloud bill that explodes at scale.
- 25 **Right-size the model to the task.** ■ Orchestrator-worker and routing patterns let a small model triage and a large model handle the hard step. *National implication:* bake cost-and-capability routing into the platform — at population scale, model-selection discipline is a budget line, not a detail.
- 26 **Engineer for latency and cost as first-order constraints.** ■ τ -bench and GAIA show agents are slow and expensive exactly where they fan out; edge/MEC and caching matter. *National implication:* set SLOs for agentic interactions and design the loop (caching, parallelisation, routing) to meet them — a query that takes minutes and costs dollars will not scale to a nation.
- 27 **Plan the model-upgrade treadmill into operations.** ■ Models are replaced every few months; an agent tuned to one model's quirks can regress on the next. *National implication:* keep an evaluation harness that re-qualifies every production agent on each model change, so upgrades are controlled rollouts, not surprises.

THE COST CLIFF

P24, P25 and P26 are one argument in three parts: agentic inference is bursty, fan-out is expensive, and naïve deployment meets a cost cliff on first contact with population scale. Managed capacity, model routing and latency SLOs are what keep the engine room affordable.

7

principles · Family C

1

■ foundational (P21)

dozens

model calls per agent task

What agents must know and remember

Everything an agent produces is a derived view, never an alternative truth. Authoritative registries are the ground truth; data readiness is a precondition, not a parallel workstream; and an agent's memory of citizens is regulated personal data.

- 28 **Authoritative registries remain the single source of truth.** ■ Agents must act on canonical systems of record, not plausible generations. *National implication:* the single most valuable thing a state can do for its agents is make population, business, land and identity registries clean, current, machine-readable and queryable.
- 29 **Data readiness is a precondition, not a parallel workstream.** ■ An agent automates whatever quality it is fed, at scale; governance, quality, lineage and interoperability come first. *National implication:* fund a data-readiness programme ahead of deployment — "agent-ready data" is the new precondition the way "machine-readable" was for open government.
- 30 **Ground generation in retrieval — RAG as the default.** ■ Pairing parametric memory with a retrieval index keeps outputs factual, specific and updatable without retraining. *National implication:* require agents on policy, law or citizen records to retrieve from authoritative sources and cite them; ungrounded generation on official matters is a defect.
- 31 **Use graph knowledge stores where relationships and verifiability matter.** ■ GraphRAG and LLM+KG enable multi-hop, verifiable reasoning that flat vector search misses — the shape of much government knowledge. *National implication:* build knowledge graphs over structured records as a first-class retrieval substrate, not only vector indexes.
- 32 **Engineer the context window deliberately.** ■ The Context Engineering survey reframes "what goes in the window" as a formal discipline of retrieval, processing and management. *National implication:* specify context assembly, compression and ordering policies — context engineering is where most agent quality and cost are won or lost.
- 33 **Exploit long context, but don't let it replace governance.** ■ Gemini 1.5's million-token recall lets large state be held in-context, but that is not data governance. *National implication:* use long context to simplify architecture where it helps, but keep authoritative retrieval and lineage — a document in context is not the source of truth.
- 34 **Make tool results and structured outputs first-class, persistable state.** ■ Function-call returns and schema-constrained outputs are auditable state that flows into memory and downstream agents. *National implication:* standardise typed schemas so agent actions are interoperable, loggable and chainable — free text between agents is a future incident.
- 35 **Govern agent memory as a regulated personal-data store.** ■ An agent's episodic and long-term memory of citizens is personal data, subject to consent, retention and deletion. *National implication:* apply data-protection law to agent memory explicitly — purpose limitation, retention limits, a right to deletion, with India's DEPA-style consent layer extended to agents.

NATIONAL IMPLICATION • GROUND TRUTH

Family D's spine is P28–P29–P35: registries are truth, readiness precedes deployment, and memory is regulated. Three of eight principles here are ■ foundational — the highest concentration of any family, because an agent that automates dirty data simply scales the error.

8

principles • Family D

3

■ foundational

1,400

papers in Context survey

No anonymous actors in production

An action you cannot attribute is an action you cannot govern. This is the densest family in the canon: identity, scoped delegation, liability, action gating, human oversight, an interruptible kill-switch and real-time observability — six of seven are foundational.

- 36 **Every agent has a verifiable identity — no anonymous actors.** Agent identifiers are the first oversight measure; an unattributable action is ungovernable. National implication: extend national identity infrastructure to issue and manage *agent* identities, so every autonomous action traces to a specific, registered agent instance.
- 37 **Authenticate and authorise agents through scoped, auditable delegation.** Authenticated Delegation extends OAuth 2.0 / OIDC with agent credentials carrying least-privilege, time-bounded scope. National implication: build "delegation infrastructure" — a citizen or civil servant granting an agent specific, revocable, logged authority — as core national plumbing.
- 38 **Establish clear liability and principal-agent accountability.** Agency theory — information asymmetry, discretionary authority, loyalty — demands someone answerable. National implication: legislate the accountability chain so every consequential action has a named responsible human or institution; "the AI did it" must never be a valid defence.
- 39 **Constrain the action space and gate high-impact actions.** Default-deny for irreversible operations; require approvals for consequential ones. National implication: configure production agents with least-privilege action sets and mandatory human approval gates on anything that moves money, changes legal status, or cannot be undone.
- 40 **Keep a human in the loop with real authority to intervene.** The one universal guardrail — the human must *review and override*, not merely watch. National implication: define mandatory oversight points and escalation paths per use-case, and ensure the reviewer has the information, time and authority to actually say no.
- 41 **Build the kill-switch and make interruptibility testable.** Attributability and interruptibility — stop and roll back — are baseline practice. National implication: require a tested emergency stop and rollback for every production agent, and rehearse it; an autonomy you cannot halt is one you cannot deploy.
- 42 **Make agents observable: monitor and log everything in real time.** Real-time monitoring and activity logging are the second and third oversight measures. National implication: capture immutable, queryable logs of every agent's reasoning, tool calls and actions — the agentic analogue of Estonia's KSI tamper-evident logging.

“

"The AI did it" must never be a valid defence. For every consequential agent action there is a named responsible human or institution, or there is no deployment.

— SYNTHESIS, NATI PRINCIPLES (FAMILY E, P38)

7

principles · Family E

6

foundational

3

oversight measures (Visibility)

How agents reach the world

Tool calls into government systems must pass through the same secure, signed, logged discipline as X-Road traffic — never direct database access. Standardise the connector, standardise agent-to-agent coordination, and treat discovery as infrastructure.

- 43 **Standardise tool connection — adopt MCP as the national "USB-C for AI."** ■ The Model Context Protocol is the emerging open standard — "build once, integrate everywhere" — the agentic counterpart to API-first and X-Road's connector. National implication: expose every registry and line-of-business system as a governed MCP server with a typed contract.
- 44 **Standardise agent-to-agent coordination — adopt A2A-style interoperability.** ■ Agent2Agent's capability-advertising "Agent Cards" let agents across vendors and ministries discover and delegate. National implication: require an open interoperability standard so a citizen-services agent can securely delegate to a tax or registry agent across agency boundaries.
- 45 **Treat tool and agent registries / discovery as infrastructure.** ■ A discoverable, governed catalogue of tools and agents is core plumbing — the analogue of Estonia's RIHA. National implication: operate a national registry of approved agents and tools with capability metadata, assurance status and access policy, so wiring is discovered and governed, not hard-coded.
- 46 **Run a control-plane orchestration runtime for long, stateful workflows.** ■ LangGraph and AutoGen supply persistence, branching, retries and human-in-the-loop across long-running work. National implication: standardise on an orchestration layer with persistence, explicit human-approval nodes and recovery — government workflows run for hours or days, not one prompt.
- 47 **Use computer-use / browser agents to bridge the legacy long tail — with extra caution.** ■ Computer use and WebArena let an agent operate human UIs directly — essential for un-API'd systems, and the highest-risk mode. National implication: permit it as a deliberate bridge under the strictest monitoring, scoping and approval — a transition, not a destination.
- 48 **Connect agents to the world through secure, governed gateways.** ■ Tool calls must pass through signed, logged exchange discipline — never unmediated database access. National implication: route all agent-to-system calls through policy-enforcing gateways (authn, authz, rate-limiting, logging), preserving the digital state's exchange model when the caller is an agent.

THE FEDERATION PRINCIPLE, REBORN

Family F is X-Road for agents. P43 (MCP) is the connector, P44 (A2A) is cross-agency federation, P45 is the RIHA-style registry, and P48 is the signed gateway. The digital state's interoperability canon is not replaced — it is extended to a caller that reasons and acts.

6

principles · Family F

2

■ foundational (P43, P48)

MCP

"USB-C for AI"

Governance is the load-bearing wall

The smallest family carries the two most absolute rules: no deployment without an evaluation harness, and a binding national risk-management spine. For agents that *act*, governance is not paperwork — it is architecture.

- 49 No deployment without an evaluation harness — and continuous re-evaluation.** ■ AgentBench (competence), GAIA (real-world assistance) and τ -bench (policy adherence, the pass^k reliability metric) measure capability *and* consistency; τ -bench's repeated-trial failures are the single most important deployment caution in the collection. *National implication:* require every production agent to pass a documented eval suite — capability, reliability under repetition, safety, policy-adherence — re-qualified continuously and on every model change.
- 50 Adopt a national risk-management spine and a trustworthy-AI bar.** ■ NIST AI RMF (Govern · Map · Measure · Manage) with its GenAI Profile, the EU AI Act's risk tiers, and Singapore's Model AI Governance Framework with AI Verify define the master control structure — validity, safety, security, accountability, transparency, explainability, privacy and fairness. *National implication:* adopt one framework as the binding national spine, map every system to its risk tier, and make conformity a precondition of deployment.

THE DEPLOYMENT CAUTION

τ -bench's finding that agents frequently fail on *repeated-trial consistency* — passing once, then failing the same task on retry — is why P49 demands reliability-under-repetition, not just a one-shot benchmark. An agent that works in the demo and fails the fourth caller is not deployable.

How the 50 compose: the agentic build sequence

The principles are a dependency graph, read as a sequence of four tiers.



“

The agentic state is not a leap past the digital state — it is the digital state's canon (identity, registries, federation, interoperability, governance, oversight) extended to a worker that reasons and acts.

— **SYNTHESIS**, NATI PRINCIPLES (THE SPINE)

2

principles · Family G

2

■ foundational

50 / 50

principles covered

The 50 Components — A Bill of Materials for the Agentic State

Seven layers stack from the engine room to the governance spine: compute & models, data & knowledge, memory & state, agent runtime, identity & trust, tools & integration, and governance & people. Each component carries a priority tier — P0 Bedrock through P4 Frontier — and an honest count of how many nations actually have it. Read down the tiers and a single finding emerges: the leaderboard is short, and it gets shorter fast.

L1 Compute & Model Substrate — The Engine Room

An agent is only as capable as the model it runs on and the silicon that serves it. Layer 1 is the engine room: governed access to frontier reasoning, a sovereign fallback for data that cannot leave the country, and the burst capacity that population-scale agent fan-out devours. Two of its seven components are P0 — without model access and serving capacity there are simply no agents.

#	COMPONENT	TIER	WHAT IT IS	ADOPTION
1	Frontier model access (governed)	P0	Assured, governed access to frontier LLMs — the reasoning engine every agent runs on	~15–25 nations (est.)
2	Sovereign / open-weight inference path	P1	In-country, controllable inference for data that cannot leave national control	~10–15 nations (est.)
3	Managed inference capacity	P0	GPU/accelerator serving capacity for inference-heavy, bursty agent workloads	~20–30 nations (est.)
4	Model gateway / abstraction layer	P1	A stable internal API hiding specific providers — optionality, no lock-in	~10–20 nations / inst. (est.)
5	Cost-and-capability routing	P2	Sends each step to the cheapest model that can do it	Few national platforms (est.)
6	Inference observability & FinOps	P2	Telemetry on token spend, latency, throughput per agent, with quotas	Minority of platforms (est.)
7	Edge / low-latency inference	P3	Inference pushed to the edge for latency-sensitive, data-local tasks	Handful of nations (est.)


DEEP DIVE • #1 FRONTIER MODEL ACCESS • P0

Effectively every nation can *buy* API access; almost none has **governed** access — central contracts plus a routing layer so agencies draw only on vetted models. The flagships are the US **USAi** platform and the EU's **GenAI4EU**; governed national programmes exist in an estimated 15–25 nations, clustered among the US, UK, Singapore, the UAE and the Gulf, and China. The gap between "can buy" and "governs" is the first place the curve bends.

DEEP DIVE • #3 MANAGED INFERENCE CAPACITY • P0

Agent tasks fan out into many model calls, so serving capacity — not model quality — is the binding constraint at scale. National AI-compute initiatives (the US AI Action Plan infrastructure pillar, EU, UK, Gulf, India, Japan, Korea) put roughly 20–30 nations in the frame, but explicit capacity reservation tuned for bursty agentic load remains rarer than the headline count suggests.

FIG 15.1



Layer 1 adoption — approximate nation-counts (est.). Midpoints of the synthesis bands; all figures order-of-magnitude, not a census. The engine room is the most widely held layer — and even here the governed, routed variants thin out fast.

L2 Data @ Knowledge Substrate — What Agents Know

Agents automate whatever data quality they are fed; bad registries plus agents produce errors at machine speed. Layer 2 is the substrate of ground truth — authoritative registries, a data-readiness programme, and the retrieval pipeline that grounds every official answer in a citable source. Three of its nine components are P0: registries, governance, and RAG are not parallel work, they are prerequisites.

#	COMPONENT	TIER	WHAT IT IS	ADOPTION
8	Authoritative, agent-ready registries	P0	Clean, current, machine-readable systems of record — the ground truth agents act on	~40 base; fewer agent-ready (est.)
9	Data-readiness & governance programme	P0	Quality, lineage, access rights, stewardship over data agents consume	~30–50 nations (est.)
10	Retrieval-Augmented Generation (RAG) layer	P0	Pipeline grounding agent output in authoritative sources	Most serious pilots (est.)
11	Vector knowledge store	P1	Dense embedding index for semantic search over unstructured corpora	Standard where RAG runs (est.)
12	Graph knowledge store / GraphRAG	P2	Entity-relationship graphs for multi-hop, verifiable reasoning	Minority of platforms (est.)
13	Context-engineering layer	P1	Discipline of assembling, compressing, ordering what enters the context window	Formalised in few (est.)
14	Long-context capability	P2	Million-token models holding large state in-context	Any frontier-model user
15	Structured-output / schema layer	P1	Typed, schema-constrained outputs that are auditable and chainable	Standard in mature builds (est.)
16	Consent / data-empowerment for agents	P2	Consent infrastructure governing what data agents may access & remember	DEPA-style ~15; agent-aware few (est.)


DEEP DIVE • #8 AGENT-READY REGISTRIES • P0

Interoperable digital base registries exist in roughly 40 nations, but *agent-ready* — clean, exposed via governed APIs, with lineage an agent can cite — is a higher bar that fewer clear. This is the inheritance thesis in miniature: the nations that built the digital foundation start the agentic race with this component half-built; the nations that did not will find that buying agents on dirty registries yields confident, fast, wrong answers.

DEEP DIVE • #10 RAG LAYER • P0

Ungrounded generation on official matters is a defect, not a feature. The RAG pipeline — chunk, embed, index, rewrite, re-rank — wires an embedding model and a vector index into the agent loop so every answer traces to an authoritative source. It is ubiquitous in private deployments and a standard component of most serious national pilots; the differentiator is no longer *having* RAG but governing its corpus and provenance.

FIG 15.2



Layer 2 adoption — approximate nation-counts (est.). Inherited digital-state counts (registries, governance, consent) are widely held; the distinctively agentic refinements — agent-ready data, agent-aware consent — collapse toward single digits.

L3 Memory & State — What Agents Remember

Memory turns an agent from a one-shot tool into an institution that carries a case across sessions and learns from experience. Layer 3 climbs from the context-window scratchpad every agent has, through OS-style paging and episodic streams, to the governed memory store the law actually requires — because an agent's memory of citizens is personal data. Two P0s anchor it: working memory, and the regulated store that makes the rest lawful.

#	COMPONENT	TIER	WHAT IT IS	ADOPTION
17	Short-term / working memory	P0	The context window as the agent's active scratchpad for the current task	Universal to any agent
18	Long-term memory with paging	P1	OS-style tiered memory (MemGPT) for an unbounded effective horizon	Leading platforms (est.)
19	Episodic memory	P2	Time-stamped memory stream retrieved by recency × importance × relevance	A minority (est.)
20	Semantic memory & reflection	P2	Distilled higher-level facts synthesised from raw episodes, with forgetting	Research→early prod. (est.)
21	Governed memory store (regulated data)	P0	Agent memory treated as regulated personal data — purpose limits, retention, deletion	Operationalised in very few (est.)
22	Skill / workflow library	P2	Catalogue of validated, reusable agent skills shared across agencies	A few leaders (est.)


DEEP DIVE • #21 GOVERNED MEMORY STORE • P0

The moment an agent remembers a citizen, its memory becomes a regulated personal-data store — subject to purpose limitation, retention limits and deletion rights, with audit over every read and write. The governance literature recognises this as a *requirement*; almost no nation has operationalised it. It is the quietest P0 in the catalogue and one of the most consequential: an ungoverned memory is a standing data-protection breach that compounds with every interaction.

DEEP DIVE • #22 SKILL / WORKFLOW LIBRARY • P2

A versioned national repository of approved agent workflows is the agentic eGovFrame — build once, reuse many. Following Voyager's skill-library pattern, validated skills become shared assets across ministries rather than re-implemented per agency. It is nascent: a few leaders are assembling catalogues, but most builds still reinvent the same casework loop in every department, paying the integration cost repeatedly.

FIG 15.3



Layer 3 adoption — approximate nation-counts (est.). Working memory is free with every agent; the governed, regulated memory store the law demands is the rarest piece on the page — a P0 honoured in principle and skipped in practice.

L4 Agent Runtime & Patterns — How Agents Work

Layer 4 is the mechanics of agency: the reason-act-verify loop, the tool-calling that turns a text generator into a worker, the approval gates that hold a human in command of consequential action, and the orchestration runtime that carries a workflow over hours or days. Four P0s define it — the loop, tool use, human-in-the-loop gates — and they are where "demo" becomes "infrastructure."

#	COMPONENT	TIER	WHAT IT IS	ADOPTION
23	Standard agent loop (reason-act-verify)	P0	Explicit, inspectable gather→act→verify loop — the atomic unit of agency	De-facto standard (est.)
24	Reflection / self-correction	P1	Generate-critique-revise loops that catch the system's own errors	Leading deployments (est.)
25	Planning & task decomposition	P1	Decompose a goal into an ordered, adaptable plan exposed for audit	Standard in complex agents (est.)
26	Tool-use / function-calling	P0	Deciding when/how to call external tools and folding results back	Universal; safe use is the gap
27	Orchestration runtime (stateful)	P1	Control plane for long-running, multi-step workflows with persistence & recovery	Leading platforms (est.)
28	Multi-agent orchestration	P2	Manager agent decomposing work and delegating to specialised workers	Front-runners only (est.)
29	Human-in-the-loop approval gates	P0	Mandatory review/override before consequential or irreversible actions	8+ US states; leaders (NASCIO 2026)
30	Verification / self-test step	P1	A check or second agent before committing an action	Leading practice (est.)
31	Computer-use / browser-agent	P3	Agents operating human UIs directly for un-API'd legacy systems	Pilots in a few nations (est.)
32	Workflow-vs-agent decision standard	P2	Published guide mapping task properties to the simplest sufficient pattern	Rare as national policy (est.)


DEEP DIVE · #29 HUMAN-IN-THE-LOOP GATES · P0

The universal guardrail across the UK Playbook, NASCIO and the EU AI Act: mandatory review and override before any irreversible action, implemented as approval nodes in the runtime. One of the few agentic P0s with a **hard signal** — 8+ US states run agentic tools with these gates operationalised (NASCIO 2026).

DEEP DIVE · #23 STANDARD AGENT LOOP · P0

An explicit gather-context → act → verify loop (ReAct, the Claude Agent SDK) is the atomic unit of agency and the basis of interpretability: an inspectable loop is an auditable agent. It is de-facto standard in every serious build — so the differentiators live upstairs in orchestration and downstairs in trust.

FIG 15.4



Layer 4 adoption — approximate counts. #29 carries a NASCIO 2026 hard signal (8+ US states); the rest are est. midpoints. The loop and tool-use are universal; governed, multi-agent variants are front-runner territory.

L5 Identity, Trust & Control for Agents

This is the load-bearing layer of the agentic state — the agentic extension of eID. An unattributable action is an ungovernable one, so every agent needs a verifiable identity, a scoped and revocable mandate, a least-privilege action space, an immutable log, and a named human on the legal hook. Six of its eight components are P0, more than any other layer, and they are exactly where adoption is thinnest.

#	COMPONENT	TIER	WHAT IT IS	ADOPTION
33	Agent identity service	P0	Verifiable identities for every agent instance — attribution of action	Single digits of nations (est.)
34	Authenticated delegation infrastructure	P0	Scoped, time-bounded, revocable, logged authority — OAuth/OIDC for agents	Very few nations (est.)
35	Action-space constraint / least-privilege	P0	Per-agent action sets, default-deny on irreversible operations	Inconsistently done (est.)
36	Kill-switch & interruptibility	P1	A tested emergency stop and rollback for every production agent	Tested in few (est.)
37	Agent observability & immutable logging	P0	Real-time monitoring + immutable logs of reasoning, tool calls, actions	Leaders (est.)
38	Liability & accountability framework	P0	Legal chain naming a responsible human for every consequential action	Binding regimes in few (est.)
39	Guardrails & policy-enforcement layer	P1	Content/safety filters, domain rules, output provenance on agent actions	Adopted by leaders (est.)
40	Secure agent-to-system gateways	P0	Policy-enforcing gateways between agents and systems — the X-Road discipline	Inherited; agent-specific in leaders (est.)


DEEP DIVE · #33 AGENT IDENTITY SERVICE · P0

Issuing and managing verifiable identities for every agent instance is the foundation of the whole trust layer — an action no one can attribute is an action no one can govern. The fix is to extend existing national identity infrastructure to issue agent identities. It is a recognised requirement, yet production-grade national agent-identity exists in only single digits of nations. This is the single steepest fall in the catalogue.

DEEP DIVE · #34 AUTHENTICATED DELEGATION · P0

The agentic extension of eID: the ability for a citizen or civil servant to grant an agent a scoped, time-bounded, revocable, logged authority — OAuth/OIDC extended for agents. It is the technical basis of "acting on behalf of," and it is emerging in only very few nations as actual plumbing. Without it, every agent action is either over-privileged or unattributable; with it, delegation becomes governable infrastructure.

FIG 15.5



Layer 5 adoption — approximate nation-counts (est.). The EU AI Act is the leading liability instrument. This is the thinnest layer in the catalogue: six P0 controls, none widely held, and the distinctively agentic ones — identity, delegation, binding liability — in low single digits.

L6·L7 Tools & Integration · Governance, Eval & People

The final two layers reach outward and upward. Layer 6 is how agents touch the world — the MCP tool layer (the national "USB-C for AI"), agent-to-agent federation, registries, and the unifying national platform every ministry should inherit. Layer 7 is the governance spine: the mandatory eval harness, the binding risk framework, the oversight institution, and the workforce that works alongside agents. Five P0s span the two.

#	COMPONENT	TIER	WHAT IT IS	ADOPTION
41	Model Context Protocol (MCP) tool layer	P0	Every system exposed as a governed MCP server — the national "USB-C for AI"	Handful of leaders (est.)
42	Agent-to-agent interoperability (A2A)	P1	Open standard letting agents discover & delegate across vendors & ministries	Nascent; pilots (est.)
43	Agent & tool registry / discovery	P1	National registry of approved agents/tools with capability & assurance metadata	Very few registries (est.)
44	Tool-calling correctness & scale tooling	P2	Map intent to correct API calls across thousands of services, no hallucination	Front-runners (est.)
45	Integration to legacy via computer-use bridge	P3	Governed computer-use as a monitored bridge to never-API'd systems	Early pilots (est.)
46	National agent platform (unifying substrate)	P0	Shared platform binding models, tools, memory, guardrails, observability	~5–15 nations (est.)
47	Agentic evaluation harness	P0	Required eval suite — capability, reliability (pass^k), safety, policy adherence	Mandatory regimes very few (est.)
48	National AI risk-management framework (binding)	P0	Master control spine — NIST AI RMF / EU AI Act tiers / Singapore MGF	Binding & applied ~10–20 (est.)
49	Oversight & assurance institution	P1	A body & process for human oversight, escalation, independent assurance	A minority of nations (est.)
50	Workforce & capability programme	P1	Training civil servants to work with agents; role-transition planning	~20–40 nations (est.)

#41 MCP TOOL LAYER · P0

Every registry and line-of-business system exposed as a governed MCP server with a typed contract — the national USB-C for AI. Platform value is dominated by the breadth of tools agents can safely call. The standard is ~18 months old; national adoption sits in a handful of leaders.


#46 NATIONAL AGENT PLATFORM · P0

The shared substrate binding models, tools, memory, guardrails and observability so every ministry inherits safety rather than re-inventing it — government-as-a-platform, agentic. Real platforms exist in ~5–15 nations: USAi, GenAI4EU, the UK, Singapore, the UAE, Estonia.

#47 EVAL HARNESS & #48 RISK FRAMEWORK · P0

τ-bench shows agents fail on consistency, so no eval means no deploy — yet mandatory regimes exist in very few nations. The binding risk framework (NIST AI RMF, EU AI Act, MGF) is the load-bearing wall; applied to agents, live in ~10–20.

FIG 15.6



Layers 6 & 7 adoption — approximate nation-counts (est.). Workforce and risk frameworks inherit broad reach; the distinctively agentic P0s — MCP, the eval harness — sit at the bottom, the gate the rest of the stack waits behind.

MOVEMENT C • §15 • THE CENTRAL FINDING

The Adoption Curve — Where Nations Fall Off

Read the 50 components down their priority tiers and the catalogue resolves into one shape: a steep descending curve. On the foundational *digital* pieces agents reuse, dozens qualify. On the distinctively *agentic* components — agent identity, delegation, eval harnesses, registries — the count with a governed, production-grade capability is single digits to low tens, falling toward one or two at the frontier.





FIG 15.7



The agentic adoption curve. Approximate nation-counts by tier (all est. except the NASCIO hard signal). The split inside P0 is the whole story: the inherited *digital* bedrock is broadly held; the *distinctively agentic* bedrock collapses to under ~15 nations and keeps falling.

STEEPER THAN THE DIGITAL STATE

For **digital** infrastructure, nearly every nation held the P0 layer and the drop-off came later, at P2–P4. For **agentic** infrastructure it starts almost immediately: the seven distinctively agentic P0 components exist *together* in fewer than ~10 nations. The curve does not taper — it cliffs.

“

The optimal agentic architecture is not exotic — it is the agentic P0–P1 components built and wired onto a healthy digital state. Almost no one has done it yet.

— NATI SYNTHESIS, B • THE 50 COMPONENTS

MOVEMENT C · §15 · THE CENTRAL FINDING

Why the Window Is Open

The short leaderboard is not a counsel of despair — it is the strategic opening. Because the agentic P0–P1 layer is held by so few, a nation that already built the digital foundation can leap to the front of a race barely two years old. The table restates the curve as a stepped ledger: how many nations clear each tier, and what the tier gates.

TIER	COMP.	NATIONS (EST.)	WHO HAS THEM & WHAT THE TIER GATES
P0 Bedrock	19	digital: dozens agentic: <~15	The <i>agentic</i> P0 — identity, delegation, eval harness, governed memory, MCP, platform, binding risk framework — is the real gate
P1 Core	16	~5–20	The "serious agentic state" band: US, the EU bloc, UK, Singapore, UAE, Estonia
P2 Differentiator	9	~3–10	GraphRAG, multi-agent orchestration, skill libraries, registries
P3 Advanced	4	~2–8	Edge inference, computer-use at scale, episodic memory in production
P4 Frontier	—	~1–3	The fully agent-native state — folded into P3

The inheritance dividend

The nations that built the digital foundation — Estonia, Singapore, the UAE, the UK, leading US states — start the agentic race with components 8 (registries), 9 (data governance), 16 (consent), 40 (secure gateways) and 46 (national platform) already half-built. The agentic state is the digital state's

canon — identity, registries, federation, interoperability, governance, oversight — extended to a worker that reasons and acts. The nations that did not face the harder lesson: **buying agents without the foundation yields demos, not infrastructure** — a fast, confident way to be wrong at population scale.

5

inherited components leaders already hold (#8, 9, 16, 40, 46)

NATI SYNTHESIS

7

distinctively agentic P0s that define the gate

NATI SYNTHESIS

METHODOLOGY NOTE

Hard figures — 8+ US states in production (NASCIO 2026), the FDA agency-wide rollout, USAi and GenAI4EU — are cited to source. Inherited digital-state counts trace to the digital components catalogue. Every (est.) figure is a reasoned, order-of-magnitude band: no global census of national agentic infrastructure yet exists, and any single precise number should be treated with caution.

The 10 Best Agentic Architecture Types

Ranked by how load-bearing each pattern is for a national-scale, accountable deployment — a blend of how decisive it is for governed autonomy, how well-evidenced it is in the field, and how sharply it separates a thinking state from a chatbot procurement. The ranking closes on the optimal reference design: the 8-layer National Agentic Technology Stack that braids all ten onto a working digital foundation.

#1 Sovereign Agent Platform

#2 Grounded RAG-Agent

#3 Orchestrator-Worker

#4 Supervised Autonomy

#5 A2A Federation

#6 MCP Tool-Augmented

#7 Stateful Workflow

#8 Reflective Agent

#9 Memory-Augmented

#10 Computer-Use Bridge

PART 1 · THE RANKED ARCHETYPES — 1 OF 6

The ten types are not rivals but emphases; a mature national stack braids several together. They are presented below to a single anatomy — **shape, why it ranks, trade-off**, exemplars and a tier/role chip — so the run reads as one comparable family.

FIG 19.1 · RANK LADDER



Load-bearing score (illustrative / est.). A blended index of how decisive each pattern is for governed national autonomy, how well-evidenced it is, and how sharply it separates a real agentic state from a chatbot. Magenta = load-bearing P0/P1 cornerstones.

01 The Sovereign Agent Platform P0 bedrock · where you build

SHAPE

A single, centrally-provisioned national platform giving every agency shared model access (behind a gateway), shared tools (via MCP), shared memory, guardrails, observability and an eval harness — at no marginal cost. Ministries build agents *on* the platform rather than buying them piecemeal.

WHY IT RANKS #1

The agentic successor to Singapore's SGTS and the UK's GDS — the single highest-leverage move a state can make. It is where every other principle is operationalised *once*: safety, identity, evaluation, cost control. Without it a nation gets a thousand ungoverned pilots; with it, governed infrastructure.

TRADE-OFF

Needs a powerful, well-funded central authority and the political capital to make the platform mandatory. Weaker fit for highly federal states unless paired with the federation pattern (#5).

- ▶ **United States (GSA)** — USAi: government-wide chat + API + evaluation console over multiple leading models, offered free to agencies; paired with OneGov procurement (folder 08).
- ▶ **European Union** — GenAI4EU / Apply AI Alliance as the bloc-scale shared-platform play.
- ▶ **Estonia / Agentic State** — the "sovereign tech stack" enabler names the same idea; GaaP lineage via OECD DGI (folder 04).

02 The Grounded RAG-Agent P0 bedrock · how you stay correct

SHAPE

Every consequential agent is a reason-act loop whose every factual claim and action is grounded in retrieval from authoritative sources — base registries, law, policy, case files — through a RAG layer (vector + graph), with citations and structured outputs. Ungrounded generation on official matters is treated as a defect.

WHY IT RANKS #2

For a state, correctness and provability outrank cleverness. The foundational RAG work and the RAG survey establish grounding as the antidote to hallucination; OECD/GPAI makes authoritative government data the ground truth agents must act on. This is what makes an agent's output *defensible* — it can show its sources.

TRADE-OFF

Quality is bounded by data readiness; on a messy registry the RAG-agent faithfully retrieves garbage. It forces the correct but expensive discipline of data-readiness first.

- ▶ **Foundational RAG** — arXiv 2005.11401 + the RAG survey 2312.10997; GraphRAG for graph-anchored retrieval.
- ▶ **OECD / GPAI** — government-data report: authoritative registries as agent ground truth (folder 11).
- ▶ **US FDA** — premarket-review agents grounded in structured submissions (folder 08).

“

For a state, correctness and provability outrank cleverness. A government that deploys ungrounded agents on official matters is automating unaccountability.

— **SYNTHESIS C**, TEN AGENTIC ARCHITECTURES & THE OPTIMAL STACK

03 The Orchestrator–Worker Multi-Agent System P1 core · how agents are built

SHAPE

A manager/orchestrator agent decomposes a complex task, delegates sub-tasks to specialised worker agents, and synthesises their results — the agentic analogue of a case team. Both Anthropic's *Building Effective Agents* and OpenAI's *Practical Guide* name this the primary multi-agent pattern.

WHY IT RANKS #3

Government work is cross-domain — one benefit claim may touch identity, tax, health and housing. Specialised agents, each with its own tools and policy scope, cooperate under one coordinator, scaling to casework a single agent cannot hold. It maps naturally onto the federated structure of government itself.

TRADE-OFF

Coordination overhead, error propagation between agents, harder debugging. Anthropic's own guidance: don't reach for multi-agent until a single agent provably cannot do the job.

- ▶ **Anthropic / OpenAI** — orchestrator-worker named as the primary multi-agent pattern in both production guides.
- ▶ **AutoGen** — arXiv 2308.08155; multi-agent surveys 2309.07864 on agent societies.

04 The Human-in-the-Loop Supervised-Autonomy Architecture

P0 bedrock · how you stay safe

SHAPE

Agents operate autonomously *within bounded action spaces*, but every consequential or irreversible action passes through a mandatory human approval gate, with escalation paths, full observability and a tested kill-switch. Autonomy is earned level-by-level up a maturity model.

WHY IT RANKS #4

The only pattern that appears in literally every credible governance source — the UK AI Playbook, NASCIO's maturity model, the EU AI Act's high-risk requirements, OpenAI's governance practices, the oversight paper (arXiv 2506.04836). For a state, supervised autonomy is not a constraint on the architecture; it *is* the architecture — what makes autonomous action politically and legally survivable.

TRADE-OFF

Human gates cap throughput and can decay into rubber-stamps when the reviewer lacks time, information or authority. The design challenge is meaningful oversight, not theatrical oversight.

- ▶ **United Kingdom** — AI Playbook (10 principles + assurance) as the canonical supervised-autonomy doctrine.
- ▶ **Leading US states** — NASCIO 2026: 8+ states run agentic tools in production under maturity-model discipline.
- ▶ **EU / OpenAI** — EU AI Act high-risk gates; OpenAI *Practices for Governing Agentic AI Systems* (folder 10/08).

05 The Federated Inter-Agent Mesh (A2A / Internet of Agents)

P1 core · how agents reach each other

SHAPE

Each agency runs its own agents over its own systems; agents discover and delegate to each other across boundaries via an open agent-to-agent standard (Agent2Agent's capability-advertising "Agent Cards"; AGNTCY's federated Agent Directory). No central monolith — a mesh of autonomous, interoperating agents, mirroring X-Road's federation at the agent layer.

WHY IT RANKS #5

The agentic reincarnation of the #1 *digital* archetype (X-Road federation). It preserves agency autonomy and data ownership while enabling cross-agency action — a citizen-services agent securely delegating a sub-task to a tax or registry agent it discovered. For large, federal, sovereignty-conscious states it is the natural shape.

TRADE-OFF

Nascent standards (A2A donated to the Linux Foundation only in 2025), immature security/trust models for cross-agent delegation, and harder end-to-end accountability across a mesh.

- ▶ **A2A → Linux Foundation** — open agent-to-agent standard; "Agent Cards" for capability advertising.
- ▶ **AGNTCY** — Internet of Agents federated directory (folder 12).
- ▶ **X-Road lineage** — the federation archetype reborn at the agent layer (folder 01).

06 The MCP-Connected Tool-Augmented State P1 core · how agents reach the world

SHAPE

Every government registry, database and line-of-business system is exposed as a governed *MCP server* with a typed contract; agents are MCP clients that integrate through one standard — "build once, integrate everywhere." The state's API estate becomes a uniform, discoverable tool-surface, behind policy-enforcing gateways.

WHY IT RANKS #6

The agent surveys and the tool-use literature (Toolformer, Gorilla, ToolLLM) are unanimous: an agent's value is dominated by the tools it can safely call. MCP is the emerging open standard that turns the whole government estate into agent-callable tools without bespoke wiring — the agentic extension of the digital state's API-first principle, and the substrate beneath every other pattern.

TRADE-OFF

Exposing systems as tools multiplies the attack surface; demands rigorous gateway-level auth, scoping, rate-limiting and logging. The standard is young and evolving.

- ▶ **Anthropic MCP** — protocol announcement + spec: the open agent-to-tool standard (folder 12).
- ▶ **Tool-use literature** — Gorilla, ToolLLM, API-Bank; Toolformer on learned tool invocation.
- ▶ **API-first lineage** — the digital state's interoperability principle, extended (folder 03).

07 The Stateful Orchestrated Workflow P1 core · durable agentic process

SHAPE

Long-running processes — a visa case, a procurement, a benefits review — are modelled as durable, stateful graphs (LangGraph-style): nodes are agents/tools/decision points, state is checkpointed and persisted, and the workflow survives failure, supports retries and interposes human-approval nodes. Process, not chat, is the unit.

WHY IT RANKS #7

Real public-service work runs for hours, days or weeks and must survive crashes, restarts and hand-offs. Orchestration runtimes provide the persistence, branching and recovery that turn a fragile prompt-chain into a durable business process — the pattern that makes agents reliable enough for casework rather than just Q&A.

TRADE-OFF

Engineering complexity; modelling government processes as graphs is real work and can ossify if not maintained.

- ▶ **LangGraph** — durable, checkpointed agent graphs with branching and recovery (folder 12).
- ▶ **AutoGen** — stateful multi-agent orchestration; the workflow-vs-agent distinction (Anthropic).

08 The Reflective / Self-Correcting Agent P2 differentiator · evaluator-optimizer

SHAPE

Agents critique and revise their own work before committing it — a generate→critique→revise loop, often with a separate evaluator agent — and store lessons from failures in episodic memory (Reflexion's verbal reinforcement learning; the evaluator-optimizer workflow). The verify stage of the Claude Agent SDK loop is the production form.

WHY IT RANKS #8

In public service, a system that catches its own mistake is worth more than one marginally more accurate that never doubts itself. Reflection (arXiv 2303.11366) measurably improves reliability without retraining and produces an auditable trail of self-correction — valuable for accountability.

TRADE-OFF

Extra latency and cost per task; reflection can entrench a confident error if the critic shares the generator's blind spot — hence prefer an *independent* evaluator.

- ▶ **Reflexion** — arXiv 2303.11366: verbal reinforcement learning from self-critique.
- ▶ **Anthropic / Voyager** — evaluator-optimizer pattern; Voyager self-verification (folder 09).
- ▶ **Claude Agent SDK** — the verify stage as the production self-correction form.

09 The Memory-Augmented Long-Horizon Agent P2 differentiator · governed memory

SHAPE

Agents carry persistent, tiered memory — working memory (context), long-term memory with OS-style paging (MemGPT), episodic memory (the recency–importance–relevance stream of Generative Agents) and consolidated semantic memory — all governed as regulated personal data. The agent carries a case across sessions and accumulates institutional knowledge.

WHY IT RANKS #9

Statelessness is the difference between a chatbot and a colleague. For recurring relationships — a citizen's multi-year benefits history, an ongoing investigation — memory is what makes the agent useful and humane. The memory literature gives the architecture; the governance literature insists it be regulated.

TRADE-OFF

Agent memory of citizens is a serious privacy and security liability — it must be purpose-limited, retention-bounded, deletable and audited, or it becomes a surveillance apparatus. The architecture is inseparable from its governance.

- ▶ **MemGPT** — OS-style paging for long-term agent memory (folder 11).
- ▶ **Generative Agents** — the recency–importance–relevance episodic memory stream.
- ▶ **Governed-memory requirement** — memory regulated as personal data (folder 10).

10 The Computer-Use Legacy Bridge P3 transition · highest-risk integration

SHAPE

For the long tail of government systems that will never expose an API, agents operate the human user interface directly — reading the screen, clicking, typing (Anthropic computer use; WebArena's web-task environment) — inside a tightly sandboxed, monitored, strictly-scoped environment, as a deliberate transition mechanism.

WHY IT RANKS #10

Every state has decades of un-API'd legacy software; without a bridge those systems are invisible to agents and the agentic state stalls at its own legacy boundary. Computer-use is the pragmatic — and genuinely powerful — bridge. It ranks last because it is also the *highest-risk* integration mode: an agent with general UI control and weak scoping is the most dangerous configuration in this entire document.

TRADE-OFF

Brittle, slow, and dangerous if under-governed; a destination only by mistake. Treat strictly as a transition while the underlying systems are API-enabled or replaced.

- ▶ **Anthropic computer use** — agents operating the GUI directly under sandboxing.
- ▶ **WebArena** — arXiv 2307.13854: realistic web-task environment for UI-driving agents (folder 12).

META-OBSERVATION · THE TEN ARE A BRAID, NOT A SHORTLIST

As with the digital archetypes, the top patterns are complementary emphases, not rivals. **#1** (sovereign platform) is *where* you build; **#4** (supervised autonomy) is *how* you stay safe; **#2** (grounded RAG) is *how* you stay correct; **#6** (MCP tools) and **#5** (A2A federation) are *how* agents reach the world and each other; **#3 / #7 / #8 / #9** are *how* the agents themselves are built; **#10** is *how* you bridge the legacy past. The optimal stack on the pages that follow is the braid of all ten.

4

P0 bedrock patterns (#1,#2,#4,+identity)

10

archetypes, all load-bearing in a mature braid

<10

nations with the P0–P1 layer actually wired

PART 2 · THE OPTIMAL NATIONAL AGENTIC STACK · 2026

The Reference Design — 8 Layers Atop the Digital State

FIG 19.2 · THE OPTIMAL 8-LAYER NATIONAL AGENTIC TECHNOLOGY STACK

L8	GOVERNANCE SPINE (cross-cutting, all layers) binding risk framework (NIST AI RMF / EU AI Act / SG Model Gov) · liability & accountability law · central agent authority · mandatory eval harness · oversight institution · workforce programme
L7	EXPERIENCE & MISSION citizen super-assistant · civil-servant copilots · policy twin · proactive life-event services — bounded, human-gated
L6	ORCHESTRATION stateful durable workflows · orchestrator-worker multi-agent · A2A inter-agency federation · human-approval nodes · recovery & retries
L5	AGENT RUNTIME reason-act-verify loop · reflection / self-correction · planning · skill library · workflow-vs-agent discipline
L4	TRUST & CONTROL (cross-cutting) agent identity · authenticated delegation (OAuth/OIDC+) · least-privilege action scopes · guardrails · tested kill-switch · observability + immutable logs · secure gateways
L3	TOOLS & INTEGRATION MCP servers over every registry/system · tool & agent registry / discovery · computer-use legacy bridge (sandboxed)
L2	DATA, KNOWLEDGE & MEMORY agent-ready authoritative registries · RAG · vector + graph stores · context engineering · governed tiered memory · consent layer · structured outputs
L1	COMPUTE & MODELS governed frontier + sovereign / open-weight models · model gateway (optionality) · managed national inference · cost routing

▼ SITS ATOP THE DIGITAL STATE ▼

DIGITAL-STATE FOUNDATION (folders 01–07) — universal broadband · eID · base registries · X-Road-style signed data exchange · sovereign cloud · zero-trust security · interoperability. *The agentic stack fails without it.*

The optimal national agentic architecture. A sovereign agent platform provisioned centrally for all of government, where every agent is a grounded, reflective reason-act-verify loop with governed memory, reaching the world through MCP and other agents through A2A, all under supervised autonomy and a binding national risk framework — sitting atop a healthy digital state. Magenta = load-bearing P0/P1 layers; L4 & L8 are cross-cutting. The win is not any single exotic pattern; it is having **all the agentic P0–P1 components wired together onto a working digital foundation** — which fewer than ~10 nations have done.

8

layers · L1 compute → L8 governance

2

cross-cutting · L4 Trust, L8 Governance

<10

nations with the full stack (est.)

Seven Load-Bearing Reasons

1 • THE PLATFORM PATTERN COMPOUNDS

Every credible leader — USAi, GenAI4EU, the *Agentic State* sovereign stack — converged on a shared platform because it lets a nation operationalise safety, identity, evaluation and cost-control once and inherit them everywhere. A thousand agency pilots cannot be governed; one platform can.

2 • GROUNDING MAKES AGENTS TRUSTWORTHY

The defining risk of LLM agents is confident error; the defining requirement of public administration is defensible decisions. RAG + authoritative registries + citations resolve the tension — the agent can always show its source. Ungrounded agents on official matters automate unaccountability.

3 • SUPERVISED AUTONOMY IS THE SURVIVABLE MODE

Unanimously across the governance sources, autonomous action on citizens must be bounded, gated, attributable, interruptible and overseen (EU AI Act, UK Playbook, NASCIO, OpenAI). This is not a brake on the architecture; it is a structural member of it — climb the maturity model, don't leap.

4 • IDENTITY + DELEGATION EXTEND THE DEEPEST PRINCIPLE

The digital state's load-bearing idea was a trusted identity at the root of every transaction. The agentic state extends it: every agent has a verifiable identity (arXiv 2401.13138) and acts only under scoped, revocable, authenticated delegation (arXiv 2501.09674) — letting an agent act *on behalf of* a citizen without dissolving accountability.

5 • MCP + A2A ARE FEDERATION REBORN

The single best digital archetype was X-Road's federation — autonomous nodes interoperating through a standard, no central monolith. MCP (agent-to-tool) and A2A (agent-to-agent) reproduce exactly that shape: integrate through open standards, preserve agency ownership, avoid a central black box.

6 • MEMORY + REFLECTION MAKE DURABLE WORKERS

Stateless, single-shot agents are toys. Tiered memory (MemGPT, Generative Agents) and self-correction (Reflexion, verify stages) make agents reliable across long-horizon casework — *provided* memory is governed as regulated personal data. Architecture and governance are inseparable.

7 • YOU EARN THE INTELLIGENCE LAYER

The *Algorithmic State Architecture* models intelligence as the capstone of DPI → Data → Interoperability → Intelligence. L1–L7 are useless without the digital state beneath them, and L7 (mission) is dangerous without L1–L6. The order is the strategy.

THE SPINE

The agentic state is the digital state's canon — identity, registries, federation, interoperability, governance, oversight — extended to a worker that reasons and acts. The win is having the agentic **P0–P1 layer wired onto a healthy digital foundation** — which fewer than ~10 nations have done.

“

The win is not any single exotic pattern; it is having all of the agentic P0–P1 components wired together onto a working digital foundation — which fewer than ~10 nations have done, and perhaps two or three have done well.

— **SYNTHESIS C**, THESIS OF THE OPTIMAL NATIONAL AGENTIC STACK

PART 2 • THE BUILD SEQUENCE

What an Optimal Nation Does, In Order

The reference design is also a sequence. Each step depends on the one beneath it; the discipline is to refuse to ship autonomy before the trust spine exists, and to roll out mission last and continuously.

- 1 Pre-requisite — fix the digital state.** Ensure eID, agent-ready authoritative registries, X-Road-style data exchange, sovereign cloud and zero-trust security are healthy (folders 01–07). Without this, stop here — no amount of agent procurement substitutes for the foundation.
- 2 L1–L2 foundation.** Stand up governed model access behind a gateway, managed national inference, the RAG + vector/graph data layer, context engineering, and governed tiered memory.
- 3 L3–L4 trust spine.** Expose systems as MCP servers behind secure gateways; build agent identity, authenticated delegation, least-privilege scopes, observability/immutable logging and a tested kill-switch. *Nothing autonomous ships before L4 exists.*
- 4 L5–L6 runtime.** Deploy the reason–act–verify loop with reflection and planning; add stateful orchestration, orchestrator-worker multi-agent, A2A federation, and human-approval nodes.
- 5 L8 governance, in parallel from day one.** Adopt a binding risk framework, stand up the eval harness and oversight institution, legislate liability, and run the workforce programme.
- 6 L7 mission, last and continuously.** Roll out citizen super-assistants, civil-servant copilots, policy twins and proactive services — starting where the time analysis says the burden is highest and the risk lowest, climbing the NASCIO maturity model as reliability is proven.

8+

US states run agentic tools in production
(NASCIO 2026)

1

platform governs what 1,000 pilots
cannot

L4

the gate before any autonomy ships

THE SPINE, RESTATED

The optimal national agentic architecture, in one sentence, is the same sentence as the digital one — extended by a single, world-changing clause: **a federated, API-first, building-block platform with a trusted identity at its root and a signed exchange layer instead of a central database — now staffed by governed, grounded, supervised agents that reason and act on top of it.** No nation has the full stack. The frontier is a composite — the US's platform, Estonia's digital foundation and framework, the UK's assurance discipline, Singapore's governance toolkit, the EU's binding risk law, the UAE's mission ambition — and the decisive fact is that the agentic race is being run on the foundations laid by the digital one.

MOVEMENT D · §22

Who's Closest — The Agentic-State Scorecard


No nation has the full optimal stack. The frontier is a composite — the US's platform, Estonia's digital foundation, the UK's assurance discipline, Singapore's governance toolkit, the EU's binding risk law, the UAE's mission ambition. The seven profiles that follow rank each on the same axis: how much of the agentic P0-P1 layer is wired onto a healthy digital state — the thing fewer than ~10 nations have done.

CASE STUDY 01 · UNITED STATES (FEDERAL)

USAi @ the sovereign agent platform

The United States is closest on the single highest-leverage move in this collection — the shared, government-wide agent platform. GSA's **USAi** gives every agency a unified chat, API and evaluation console over multiple leading models, free at the point of use; **OneGov** buys "government as one customer" with agentic deals already struck; and the FDA has taken agents from pilot to an agency-wide premarket-review rollout.

The platform pattern is the US's decisive advantage: where weaker states accumulate a thousand ungoverned pilots, Washington operationalises model access, cost-routing and an eval console once and lets agencies inherit them. The **AI Action Plan ("Winning the Race")** supplies the compute-and-infrastructure pillar beneath it. The gap is governance coherence: a fast, deliberately deregulatory rollout risks outrunning the supervised-autonomy discipline the rest of this report treats as load-bearing.



AGENTIC STACK COVERAGE – INDICATIVE READINESS (EST., 0-100)



MAIN GAP — COHERENT GOVERNANCE AT SPEED

Platform maturity is racing ahead of a unifying supervised-autonomy and oversight spine. The strategic question is whether the US can keep grounding, eval harnessing and human-gating in step with a deregulatory rollout — or automate unaccountability at machine speed. The platform is there; the governing wall must keep pace.

CASE STUDY 02 · ESTONIA

The deepest foundation — and the framework that named the era

Estonia did not start the agentic race; it defined it. The **Agentic State vision paper** (Ilves/Kilian) is the first framework for the agentic transformation of core government functions, and it could only have come from the nation with the deepest digital floor to build on: **X-Road** federation, population-scale **eID**, clean base registries and KSI tamper-evident logging.

This is the inheritance thesis made flesh. Estonia begins the agentic race with half the P0 layer already laid — a signed data-exchange bus, a trusted identity at the root of every transaction, and a RIHA registry of trusted services that is the natural ancestor of an agent registry. Its strongest archetype is therefore **federation reborn at the agent layer**: autonomous nodes interoperating through a standard, no central monolith. The gap is execution scale — moving from a world-leading framework to a production agent platform at USAi's reach.

1st

framework for agentic transformation of core gov functions (Agentic State)
AGENTIC STATE VISION PAPER 2025

X-Road

signed, federated data-exchange bus — the agentic foundation
NIIS / e-ESTONIA

~99%

e-services digital; eID + base registries agent-ready (est.)
e-ESTONIA GUIDE (est.)

Strongest archetype

#5 Federated Inter-Agent Mesh plus the
capstone intelligence layer — the agentic reincarnation of the #1 digital archetype, X-Road federation.


Mini-timeline

- ▶ 2001 — X-Road data-exchange layer goes live
- ▶ 2007 — KSI tamper-evident state logging
- ▶ 2025 — Agentic State vision paper sets the agenda
- ▶ 2026 — framework → production platform (in progress)

MAIN GAP — FRAMEWORK TO PRODUCTION

Estonia has the design and the foundation; what it must now demonstrate is the scaled, governed production platform on top of them. The enablers the Agentic State paper itself names — agent governance, a sovereign tech stack, cyber and people — are the build-out still ahead. The foundation is unmatched; the capstone is half-built.

AGENTIC STACK COVERAGE — INDICATIVE READINESS (EST., 0-100)



Assurance discipline and an evidence-led target

The UK leads on the dimension that makes autonomous action survivable: **supervised autonomy and assurance**. The **AI Playbook for the UK Government** sets ten principles plus procurement and assurance guidance for safe public-sector AI, and the Alan Turing Institute supplies the evidence base that tells the state *where* to point agents first.

Turing's **Mapping the Potential** quantified that ~41% of public-sector working time in Great Britain is spent on activities supportable by generative AI (49% in education) — exactly the "start where the time is, not where the hype is" discipline. Combined with the Playbook's human-oversight and assurance requirements, the UK is the clearest exemplar of the human-in-the-loop, earn-your-autonomy model. The gap is a unifying national agent platform at USAi scale — the UK has the governance and the targeting, but not yet the single shared substrate the leaders converge on.

10

principles in the AI Playbook +
procurement & assurance guidance
UK GDS / CABINET OFFICE 2025

41%

of GB public-sector working time AI-
supportable (49% in education)
ALAN TURING INSTITUTE 2025

GOV.UK

One Login + Service Standard — the
digital floor beneath the playbook
UK GDS (LINK-ONLY)

Strongest archetype

#4 Supervised Autonomy

The only pattern that appears in every credible governance source — and the one the UK has institutionalised earliest through the Playbook + assurance.

Mini-timeline

- ▶ 2025 — AI Playbook (10 principles) published
- ▶ 2025 — Turing "Mapping the Potential" evidence base
- ▶ 2025 — GOV.UK One Login scales the identity floor
- ▶ 2026 — national agent platform (the missing piece)

AGENTIC STACK COVERAGE — INDICATIVE READINESS (EST., 0-100)





MAIN GAP — A UNIFYING NATIONAL PLATFORM

The UK has the assurance discipline and the evidence-led targeting but lacks a single shared agent platform at USAi's reach. Governance without a platform governs nothing at scale; the next move is to wire the ten principles into one inherited substrate every department builds on.

NAIS 2.0, the nine-dimension framework and AI Verify

Singapore pairs the deepest public-service AI embedding with the world's most operational governance toolkit. **National AI Strategy 2.0** (three systems / ten enablers / fifteen actions) drives AI deep into public service, while the **Model AI Governance Framework for GenAI** sets nine governance dimensions, paired with the **AI Verify** testing toolkit that turns governance from a PDF into a test you can run.

This is the platform-plus-governance archetype: SGTS already gave Singapore a government-as-a-platform base layer and reusable whole-of-government APIs, and Singpass gives ~97% identity coverage. The combination of NAIS 2.0's embedding and AI Verify's testability makes Singapore the model for measurable, conformant agentic AI. The gap is the move from copilots — AI that assists — to bounded autonomy — AI that acts — under those same controls.



MAIN GAP — FROM COPILOTS TO BOUNDED AUTONOMY

Singapore's governance and embedding are exemplary, but its agents still mostly assist rather than act. The next move is to climb the maturity model into supervised, gated autonomy — applying the nine dimensions and AI Verify to agents that take consequential actions, not only to assistants that draft them.


AGENTIC STACK COVERAGE — INDICATIVE READINESS (EST., 0-100)




Mission ambition and a Minister of AI

The UAE leads on the experience-and-mission layer — the vision of agents as front-line public servants. The **National Strategy for AI 2031** frames agents as the "workhorses that convert policy into daily service," and the world's first **Minister of State for AI** supplies the central, top-down authority the platform pattern needs.

The UAE's strength is mission clarity and institutional will: a single accountable office, a strong platform drive, and a concrete agents-in-service vision spanning visas, logistics and citizen services, all on the UAE PASS eID, signature and SSO foundation. This is the "treat agents as a new factor of production" principle taken seriously at the top of government. The gap is the governance counterweight — open eval depth and transparent oversight to match the mission ambition, so autonomy is earned rather than simply mandated.



AGENTIC STACK COVERAGE — INDICATIVE READINESS (EST., 0-100)



MAIN GAP — OPEN GOVERNANCE & EVAL DEPTH

The mission ambition outpaces the published governance. To make top-down autonomy survivable, the UAE needs open eval harnesses, transparent oversight institutions and a binding risk framework on par with its delivery drive — so that agents that act on citizens are demonstrably grounded, gated and accountable.

Apply AI, GenAI4EU and the world's binding risk law

The EU brings the one thing no other actor has: a **binding** risk framework. The **EU AI Act (Regulation 2024/1689)** is the world's first horizontal, enforceable AI law — risk tiers, high-risk requirements, mandatory human oversight — and the **Apply AI / GenAI4EU** programmes give the bloc a sovereign-platform play with explicit agentic-AI support for public administration.

The EU's strongest archetypes are therefore **federation and governance**: the AI Act as the compliance perimeter every member state must build inside, plus A2A/AGNTCY interoperability momentum that mirrors the bloc's deep heritage in cross-border interoperability (EIF, eIDAS, Once-Only). The gap is speed — turning binding frameworks and funded programmes into deployed, production platforms before the frontier moves on. The EU writes the rules of the road faster than it builds the road.

1st

binding, horizontal AI law in the world
(risk tiers + human oversight)
EU AI ACT 2024/1689

GenAI4EU

Apply AI Alliance — sovereign platform +
agentic support for public admin
EUROPEAN COMMISSION 2025 (LINK-ONLY)

A2A

AGNTCY / Agent2Agent interoperability
momentum — federation reborn
LINUX FOUNDATION 2025

Strongest archetype

#5 Federation

+

#4 Governance

The AI Act is the binding spine (P50); A2A/AGNTCY is X-Road federation reborn at the agent layer.

Mini-timeline

- ▶ 2017 — EIF; 2014 eIDAS interoperability heritage
- ▶ 2024 — EU AI Act enters into force
- ▶ 2025 — Apply AI / GenAI4EU launched
- ▶ 2026 — frameworks → deployed platforms (the race)

AGENTIC STACK COVERAGE — INDICATIVE READINESS (EST., 0-100)



MAIN GAP — SPEED OF DEPLOYMENT

The EU's binding law and funded programmes are world-leading; its risk is turning them into production platforms slower than the frontier moves. The challenge is institutional velocity — shipping the deployed agentic infrastructure as fast as it ships the regulation.

Agents in production — the maturity-model vanguard

The most concrete agentic-production evidence in this entire collection comes not from a national government but from the US states. NASCIO's 2026 report, *Beyond Generation*, documents **8+ states already running agentic tools in production** — and crucially, they reached them by *climbing* a five-phase maturity model, not leaping.

That maturity model — assistive GenAI → context-aware GenAI → task-level automation → stateful multi-step workflows → adaptive/proactive agents — is the emblematic sequencing instrument of the whole field, and the states are its proof of concept: supervised autonomy earned level by level, with oversight at each prior level before more autonomy is granted. The strongest archetype is **supervised autonomy**. The gap is the shared platform and the identity/delegation depth that a national actor like USAi can provide but a single state cannot easily build alone.






FIG D-1



The NASCIO maturity climb. Leading US states reach production agents (P4) by ascending each phase — not by leaping to autonomy. Bar widths are illustrative of relative adoption breadth across states (est.); the 8+ in-production figure is the hard NASCIO signal.

AGENTIC STACK COVERAGE — INDICATIVE READINESS (EST., 0-100)



MAIN GAP — SHARED PLATFORM + DELEGATION DEPTH

The states have the discipline; what they lack is a shared agent platform and the identity/authenticated-delegation depth a national actor provides. Fifty separate climbs of the same maturity model is the federal case for a USAi-class substrate underneath them.

MOVEMENT D • §23 – CROSS-CUTTING

23 What the leaders share — and where states fail

Read across the seven profiles and the pattern is unmistakable. No nation has the full optimal stack, but every leader converged on the same small set of moves — and every failure traces to the same small set of omissions. The order is the strategy.

What the leaders share

A shared platform, not piecemeal

procurement. USAi, GenAI4EU and the Agentic State's sovereign stack all converged on one substrate because safety, identity, evaluation and cost control are operationalised once and inherited everywhere. A thousand agency pilots cannot be governed; one platform can.

A working digital foundation underneath.

Estonia's X-Road, Singapore's SGTS, the UK's One Login, the UAE's UAE PASS — every leader is building agents on a healthy identity-and- registry floor. The agentic race is being run on the foundations laid by the digital race.

Supervised autonomy, earned by maturity.

The UK Playbook, NASCIO's five phases, the EU AI Act and OpenAI's practices all insist autonomy be bounded, gated, attributable and interruptible — and climbed, not leapt. This is a structural member of the architecture, not a brake on it.

Grounding and an eval harness.

Singapore's AI Verify and the τ -bench/AgentBench/GAIA discipline make correctness measurable. Leaders deploy agents that can show their source and pass a documented suite before — and during — production.

THE FOUR FAILURE MODES — HOW STATES AUTOMATE DYSFUNCTION

1 • Buying agents without the foundation. Procuring autonomous tools onto broken registries, missing eID and un-API'd systems. The agent faithfully retrieves and amplifies garbage at machine speed; the result is an expensive demo, not infrastructure.

2 • Ungoverned pilots. A thousand agency experiments with no central platform — no shared guardrails, no shared observability, every ministry re-failing at safety alone.

3 • No eval harness. Deploying without measuring capability *and* reliability under repetition. τ -bench's pass^k finding — that agents fail on repeated-trial consistency — is the single most important deployment caution in this collection, and it is invisible without a harness.

4 • No kill-switch, no human gate. Autonomous action on citizens with no tested emergency stop, no rollback and no meaningful approval gate. An autonomy you cannot halt is an autonomy you cannot deploy.

“

The decisive strategic fact is that the agentic race is being run on the foundations laid by the digital race — the nations that built X-Road-class interoperability, population-scale identity and clean registries start with half the agentic P0 stack already in place.

— SYNTHESIS C, TEN AGENTIC ARCHITECTURES & THE OPTIMAL NATIONAL STACK

MOVEMENT D • §23 – THE BUILD-ORDER INSIGHT

24 The recommendations matrix

The build-order insight is the whole strategy: you earn the intelligence layer, you do not skip to it. The matrix below sequences the synthesis build order into ten concrete actions — fix the digital state, stand up L1–L2, build the L3–L4 trust spine before anything autonomous ships, then runtime, governance in parallel, and mission last and continuously.

ACTION	OWNER	BY WHEN	MEASURED HOW
0 • Fix the digital state — eID, agent-ready registries, X-Road-style exchange, sovereign cloud, zero-trust	Central digital agency	Pre-requisite	Registry quality/lineage score; eID coverage %; data-exchange uptime
1 • Stand up the national agent platform — shared model access behind a gateway, tools, memory, observability	Central agent authority	0–12 mo	Agencies onboarded; marginal cost-per-agency → 0
2 • Model gateway + managed inference — multi-model optionality, cost routing, sovereign/open-weight fallback	Platform team	0–12 mo	Models routable; cost-per-resolution; sovereign inference path live
3 • RAG + registry grounding layer — vector + graph stores, citations, structured outputs	Data office	6–18 mo	% of agent claims grounded & cited; hallucination defect rate
4 • Expose systems as MCP servers behind policy-enforcing gateways; tool & agent registry	Platform + agencies	6–24 mo	Registries exposed as governed tools; gateway-mediated calls %
5 • Agent identity + authenticated delegation — scoped, revocable, logged; nothing autonomous ships before this	Identity authority	6–18 mo	Agents with verifiable identity %; delegation scopes audited
6 • Observability, kill-switch & human gates — immutable logs, tested emergency stop, approval nodes	Central agent authority	6–18 mo	Kill-switch drill pass; % consequential actions gated; log coverage
7 • Eval harness, mandatory — capability + pass [^] k reliability + safety + policy-adherence; re-qualify on every model change	Oversight institution	0–18 mo, continuous	% production agents passing documented suite; re-qual cadence
8 • Binding risk framework + liability law — NIST AI RMF / EU AI Act / SG Model Gov; named responsible human per action	Legislature + central authority	0–24 mo, parallel	Every system mapped to a risk tier; accountability chain legislated
9 • Roll out mission services — citizen super-assistant, civil-servant copilots, proactive life-events — highest-burden, lowest-risk first	Line agencies	12 mo+, continuous	Citizen outcomes resolved; civil-servant hours returned; maturity-phase

THE BUILD-ORDER INSIGHT

Rows 0–8 are the trust spine; row 9 is the only one citizens see — and it ships last, on top of everything above it. The Algorithmic State Architecture models intelligence as the capstone of DPI → Data → Interoperability → Intelligence. L1–L7 are useless without the digital state beneath them, and L7 (mission) is dangerous without L1–L6. **The order is the strategy.**

MOVEMENT E • §24 – RISKS & COUNTER-ARGUMENTS

25 Could the thesis be wrong?

A confident report owes its reader the strongest case against itself. The thesis — that nations should wire the agentic P0–P1 layer onto a healthy digital foundation now — faces four serious objections, each grounded in this collection's own evidence. None overturns the spine; each sharpens it.

OBJECTION 1 • AGENT RELIABILITY

The pass^k problem. τ -bench's central finding is that agents which pass a task once often fail it on repeated trials — consistency, not peak capability, is the binding constraint. GAIA reports models at ~15% on real-world assistance where humans score ~92%. A state that deploys agents on consequential casework before reliability is proven is automating an unaccountable error rate.

Rebuttal: this is precisely why the eval harness (P49) and supervised autonomy (P40) are foundational, not optional — the thesis *requires* measuring pass^k and gating on it, which is the opposite of naïve deployment.

OBJECTION 2 • HYPE CYCLE

The standards are two years old. A2A was donated to the Linux Foundation only in 2025; MCP is young and evolving; the agentic archetypes are mostly proven in the private sector and in pilots, not as governed national infrastructure. Building national plumbing on moving standards risks expensive churn. **Rebuttal:** the load-bearing investments — clean registries, eID, delegation, observability, an eval harness — are standard-agnostic. They pay off regardless of which agent protocol wins, which is why the build sequence front-loads them and treats the protocols as the replaceable top layer.

OBJECTION 3 • SAFETY INCIDENTS

Agent memory is a surveillance apparatus in waiting. Long-horizon agents that remember citizens, and computer-use agents with general UI control, are the two highest-risk configurations in this report. A single ungoverned incident could set public trust back a decade. **Rebuttal:** the architecture is inseparable from its governance — governed tiered memory as regulated personal data (P35), least-privilege scopes, kill-switch and immutable logs are not add-ons but structural members. The thesis bans the dangerous configuration, it does not enable it.

OBJECTION 4 • THE OECD "LIMITED PRECEDENT" CAUTION


The OECD's survey of 200 government AI use cases found the public sector has "limited precedent for autonomous software in government" while the private sector races ahead. This is the steelman's strongest card: government has never run software that *acts* on citizens at scale, and caution is warranted. But the OECD frames it as *both* a warning *and* an opportunity — the precedent gap is exactly why building the governed foundation now, deliberately and in sequence, is the responsible path rather than waiting for the private-sector pattern to arrive ungoverned.

MOVEMENT E · §24 – OUTLOOK

26 Base, bull and bear to 2030

Where does national agentic adoption land by ~2030? The honest answer is a fan of scenarios, anchored to the one hard signal we have — fewer than ~10 nations have wired the agentic P0–P1 layer onto a healthy digital foundation today.

FIG E-1



Nations with a governed agentic P0–P1 layer on a healthy digital state. Reasoned scenario estimates (est.) from the <10-nation 2026 baseline. Spread is governance velocity, not model capability.

SCENARIO	WHAT DRIVES IT	WHAT IT LOOKS LIKE IN 2030
BEAR	Frameworks stall as PDFs; pilots churn; a safety incident freezes trust; standards fragment	~15 nations; agents stuck at assistive copilots; autonomy politically radioactive
BASE	The platform pattern proves out; leaders scale to production; fast-followers inherit USAi/AI-Act-class blueprints	~30 nations with governed P0–P1; supervised autonomy normal in high-volume, low-risk functions
BULL	MCP/A2A standards diffuse; shared platforms + open eval harnesses become reusable public goods	~50 nations; the agentic capstone becomes the new digital-government baseline, as e-services did after 2010

<10	198	~38	8+
nations with the layer wired today (est.)	economies in the World Bank GTMI	digital governments OECD-benchmarked	US states already in production (NASCIO)

THE THROUGH-LINE OF EVERY SCENARIO

The variable that separates bear from bull is not model capability — it is governance velocity and whether nations build the foundation first. The states that built X-Road-class interoperability and population-scale identity start every scenario with half the agentic P0 stack in place; the rest find no amount of agent procurement substitutes for it.

MOVEMENT E • §25 — METHODOLOGY

27 How this evidence base was built

This report rests on a curated, verified corpus assembled with a single discipline: every URL was checked to resolve, only real multi-page PDFs were downloaded, and living specifications were recorded as links, not faked as downloads — an evidence base you can audit, not a bibliography you must trust.

12

thematic folders (7 digital + 5 agentic)

12

parallel research teams (7 + 5), one per folder

142

real PDFs downloaded & checked (96 + 46)

~174

sources catalogued in total (111 + 63)

The four-part discipline

Verify every URL. Each source was found via live web search and its URL verified to resolve before it entered the catalogue. Nothing was cited from memory or assumed to exist.

Only real PDFs. Where a free PDF existed, it was downloaded and checked with `file / pdftinfo` to confirm it was a real, multi-page document. 142 sources passed this bar (96 digital + 46 agentic).

Link-only for living specs. Canonical specs published as living HTML or Git (X-Road ARC-

G, GovStack, MOSIP, MCP, A2A, AGNTCY) and bot-blocked pages (OECD, UAE, EU, GSA) were recorded as **LINK-ONLY** with a working URL — never fabricated as downloads.

Grounding & est. discipline. Every number traces to the folder-04 indices or a named source. Agentic infrastructure is ~2 years old with no ITU-grade census, so most adoption figures carry an explicit **(est.)** label; only hard signals — 8+ US states (NASCIO), the FDA rollout, USAi/GenAI4EU — are cited as fact.

100%

of URLs verified to resolve before cataloguing

INTEGRITY NOTE

~45

sources on actual national architectures (~40% — above the one-third bar)

00_README

0

fabricated downloads — living specs marked LINK-ONLY, never faked

METHODOLOGY

WHY THE DISCIPLINE MATTERS FOR AN AGENTIC REPORT

A report arguing that states must ground their agents in verifiable, cited, authoritative sources cannot itself rest on unverified claims. The grounding rule we recommend for national agents — show your source, label your estimates — is the same rule this evidence base was built under. Method is the argument.

MOVEMENT E · §26 – THE EVIDENCE BASE

28 The corpus — 12 folders, ~174 sources

The full collection is twelve thematic folders — the seven that document the digital state and the five that document its agentic capstone. The counts below are the catalogued totals; together they form the spine of every claim in this report.

142

real PDFs downloaded & verified
(96 digital + 46 agentic)

00_README + AGENTIC

~174

total sources catalogued (111
digital + 63 agentic)

00_README + AGENTIC

~45

sources on real national
architectures (~40% of the
digital base)

00_README

#	FOLDER	SCOPE	SOURCES
01	Architectures · Europe	Estonia X-Road, Finland, Denmark, Norway, NL, Germany — real stacks	16
02	Architectures · Asia	Singapore, India Stack, S. Korea, Japan, China, UAE — real stacks	16
03	EU/UK/US frameworks	EIF, eIDAS/EUDI, NIST, FedRAMP, FEAR, GAIA-X, Zero Trust	15
04	Indices & frameworks	ITU, UN, World Bank, OECD, GSMA, NRI — the measurement layer	15
05	Digital Public Infrastructure	UNDP/G20, GovStack, MOSIP, Mojaloop, X-Road, DEPA	15
06	Technical architecture	Broadband, 5G/6G, Open-RAN, IXP, RPKI, DNS, edge, sovereign cloud, ZT	18
07	Emerging-nation reference	Smart Africa, Diia, PIX, Rwanda, ASA, GEA reference models	16
08	Agentic strategies · national	Agentic State, NASCIO, UK Playbook, NAIS 2.0, USAi, AI Action Plan, OECD, TBI, Turing, UAE, EU	13
09	Agent architectures & patterns	ReAct, Reflexion, Toolformer, ToT, CoT, surveys, Generative Agents, Voyager	13
10	Requirements & foundations	NIST AI RMF, EU AI Act, Visibility, OpenAI practices, AgentBench/GAIA/t-bench, SG Model Gov	11
11	Agent data & state	RAG, MemGPT, Generative Agents, GraphRAG, KG+LLM, context engineering, OECD govt-data	10
12	Agent tools & integration	MCP, A2A, Gorilla, ToolLLM, API-Bank, WebArena, computer-use, AutoGen, LangGraph, AGNTCY	12

7 + 5

folders: digital state + agentic capstone

~45

sources on real agent architectures
(folders 08–12)

~17

canonical specs kept LINK-ONLY, never
faked

MOVEMENT E · §26 – THE EVIDENCE BASE

29 The marquee sources

From ~174 catalogued sources, these ~14 are the load-bearing references — the documents the principles, components and architecture rankings most directly rest on. Each is a real, verified source or a canonical LINK-ONLY specification.

Governance, risk & identity

- ▶ **NIST AI RMF** — the risk spine (100-1 + 600-1)
- ▶ **EU AI Act** — world's binding risk law (2024/1689)

National strategies & adoption

- ▶ **Agentic State** — first agentic-gov framework (Ilves/Kilian)
- ▶ **USAi** — GSA government-wide platform
- ▶ **NASCIO 2026** — 5-phase model; 8+ states in production
- ▶ **UK AI Playbook** — 10 principles + assurance

Agent architecture & tools

- ▶ **ReAct** — the reason-act-observe loop
- ▶ **MCP** — Model Context Protocol, "USB-C for AI"
- ▶ **MemGPT** — OS-style memory paging
- ▶ **A2A** — Agent2Agent → Linux Foundation

Digital-state foundation

- ▶ **X-Road** — NIS federation; the #1 archetype
- ▶ **ITU Facts & Figures 2024** — 5.5bn online / 96% coverage
- ▶ **India Stack** — Aadhaar / UPI / DEPA
- ▶ **World Bank GTMI** — 198 economies on GovTech maturity

142

real PDFs

~174

sources total

12

thematic folders

FROM CATALOGUE TO CANON

These marquee sources are why the report can make its central claim with confidence: the agentic state is the digital state's canon — identity, registries, federation, interoperability, governance, oversight — extended to a worker that reasons and acts. Every one of those words traces to a verified document above.

The order is the strategy

The agentic state is not a leap past the digital state — it is the digital state's canon, identity, registries, federation, interoperability, governance and oversight, extended to a worker that reasons and acts. An agent is only as good as the identity it can verify, the registry it can trust, the tool it can call, and the human who can stop it. The win is not any single exotic pattern; it is having the agentic P0–P1 layer wired onto a healthy digital foundation — which fewer than ~10 nations have done, and perhaps two or three have done well. The nations that built X-Road-class interoperability, population-scale identity and clean registries start the agentic race with half the foundation already laid; the rest will find that no amount of agent procurement substitutes for it. Build the foundation. Earn the capstone. Govern the worker that acts.

[NATI principles](#)

[agentic components](#)

[architecture archetypes](#)

[verified sources](#)

ENSI · [European National Strategy & Intelligence Report](#) · [National Agentic Technology Infrastructure](#) Cite as · [ENSI, NATI: The Agentic State \(June 2026\)](#) Contact · [intelligencestrategy.org](#)