Digital Government Authority
هيئة الحكومة الرقمية

# AI Agents as Government Partners

07 December, 2025
Document Type: Brief Study
Document Classification: Public
Issue No.: 1.0

# Contents

# 1. Introduction

Governments today face dual pressures: internally, they must coordinate decisions that span multiple ministries and datasets; externally, they must respond quickly to rising citizen expectations for accessible, consistent service. Yet in both areas, existing systems fall short. Critical decisions, such as approving business permits, allocating benefits, or managing crises, often require inputs from siloed entities, slowing response times and creating blind spots. At the same time, service delivery channels like hotlines and online portals struggle with high volumes, leading to delays, inconsistent answers, and staff burnout.

AI agents offer a new digital capability for government: autonomous, always-on systems that can reason across internal silos, communicate with citizens, and support public servants in real time. Orchestration-grade agents can access data across ministries without moving it, analyze it intelligently, and generate a single, explainable recommendation. This compresses multi-agency decisions from weeks to hours. Meanwhile, conversational agents act as first-line digital clerks, resolving citizen queries instantly in Arabic and English and escalating only complex cases to human staff. Together, these AI agents unlock faster decisions, more responsive services, and a more efficient public sector.

The case for AI agents is reinforced by the latest market and performance data. The AI Index 2025 from Stanford shows AI benchmark scores rising sharply in 2024, with improvements by 49 points on math benchmark performance and 67 points on software engineering benchmarks. In addition, adoption accelerated, with 78 percent of organizations globally now using AI, up from 55 percent a year earlier. Meanwhile, global investment momentum is unprecedented: Gartner projects $644 billion in AI spending in 2025 (up 76 percent from 2024), IDC forecasts a 31.9 percent compound annual growth rate through 2029, reaching $1.3 trillion, and UBS estimates $375 billion in AI infrastructure spending in 2025, rising to $500 billion in 2026. Together, these figures underline that AI is no longer experimental; it is a strategic foundation for public and private institutions alike.

Looking ahead, new frontiers such as multi-agent systems are emerging. These are environments where several AI agents collaborate, coordinate, and divide complex tasks. Early open-source frameworks like CrewAI demonstrate how multiple agents can work together in research, planning, and service delivery. For decision makers, this points to a future where governments can deploy entire teams of agents, working seamlessly to improve efficiency, foresight, and resilience across public administration.

In Saudi Arabia, AI agents directly support the Kingdom's Vision 2030 goals of a seamless, integrated digital government. Platforms like Absher, Tawakkalna, and Najiz already provide digital access to services, but AI agents can take the next step by reasoning across ministries, delivering instant citizen support, and enabling data-driven policy decisions. Backed by initiatives from the Digital Government Authority (DGA) and the Saudi Data and AI Authority (SDAIA), these agents can help ministries accelerate approvals, ensure consistent answers, and respond faster to public needs, all while keeping full control over sensitive data.

Despite the potential of AI agents, their adoption in government raises important questions around trust, accountability, and fairness. Who is responsible when an AI agent makes a policy recommendation or guides a citizen through a sensitive process? How can entities ensure that AI decisions remain explainable, auditable, and free from bias, especially in areas like benefits eligibility screening? These systems must also navigate the linguistic complexity of serving citizens in Arabic across diverse regional dialects and formal administrative language. Without thoughtful design, an AI agent may misunderstand a citizen query or deliver a response that feels impersonal or inaccurate. Addressing these risks requires strong safeguards, including human oversight, legal clarity, and transparent decision-making methodology. Without this foundation, even the most advanced AI agent risks eroding public confidence.

This study serves as a brief, descriptive, and exploratory guide, aiming to introduce AI agents and examine their potential role in government. The sections that follow offer a roadmap for deploying AI agents in government, beginning with a definition of the technology and its evolution, followed by a framework for how AI agents can work in practice. The study then turns to real-world examples from both global and Saudi contexts, before addressing challenges like the deployment of AI Agents. It closes with recommendations to help public-sector leaders unlock the benefits of AI agents while managing their risks.

## Faster Decisions, Better Services, and Trusted AI.

Modern governments face the challenge of coordinating across ministries and respond to citizens faster than ever, but current systems are too slow, fragmented, and manual. AI agents help by reasoning across entities and providing instant, explainable support to both citizens and staff. They make decisions faster and services smarter, without moving sensitive data.

# 2. Definition of AI Agents

AI agents are autonomous or semiautonomous software entities that perceive their environment, make decisions, and act to achieve predefined goals. According to Gartner, an AI agent "uses AI techniques to perceive, decide, and act in its digital or physical environment, adapting its behavior over time."

For governments, this means software that can ingest real-time data (laws, sensor feeds, citizen requests), reason over policy rules, and execute or recommend actions, while learning from every interaction. AI agents extend traditional automation by adding continuous learning and goal-directed behavior, making them suitable for dynamic public-sector contexts where rules and conditions change frequently.

**Gartner's capability model for agents can be distilled into four practical steps for public administration:**

**1**

### Perceive

Collect data from internal systems, sensors, and external sources.

**2**

### Decide

Apply machine learning, business rules, and policy constraints to select the best action.

**3**

### Act

Execute tasks or surface recommendations through secure government platforms.

**4**

### Learn

Automatically capture feedback, update internal models, and refine future decisions.

## Core Levels of AI Agents

### Basic Agent

Focused automation

Handles predictable, rule-based tasks in stable settings (e.g., FAQ chatbots, form validation).

- **Policy fit:** Low-risk services.
- **Cost/Risk:** Minimal development and oversight. Ideal entry point for pilots.

### Intermediate Agent

Context-aware assistant

Operates across multiple data streams, learns from feedback, and optimizes trade-offs.

- **Policy fit:** Moderate-complexity decisions that still require.
- **Cost/Risk:** Higher build cost; delivers bigger efficiency gains.


### Advanced Agent

Autonomous decision partner

Functions in complex, uncertain environments, collaborating with humans and other agents.

- **Policy fit:** High-stakes, dynamic domains with strong.
- **Cost/Risk:** Maximum value but demands rigorous testing, oversight, and security.

While Intermediate Agents deliver larger efficiency gains, they operate across multiple data streams and optimize trade-offs, causing them to cost more to build than Basic Agents. Governments must weigh costs and benefits carefully before scaling Intermediate Agents beyond pilot programs.

# 3. History of AI Agents as Government Partners

The history of AI Agents as Government Partners traces how software evolved from deterministic expert systems to autonomous learning entities capable of operating in complex public-sector environments. The timeline that follows highlights the pivotal breakthroughs, symbolic reasoning, speech understanding, autonomous planning, deep-learning, and large-scale language models, that gradually equipped agents to perceive, decide, and act alongside human actors in real time.

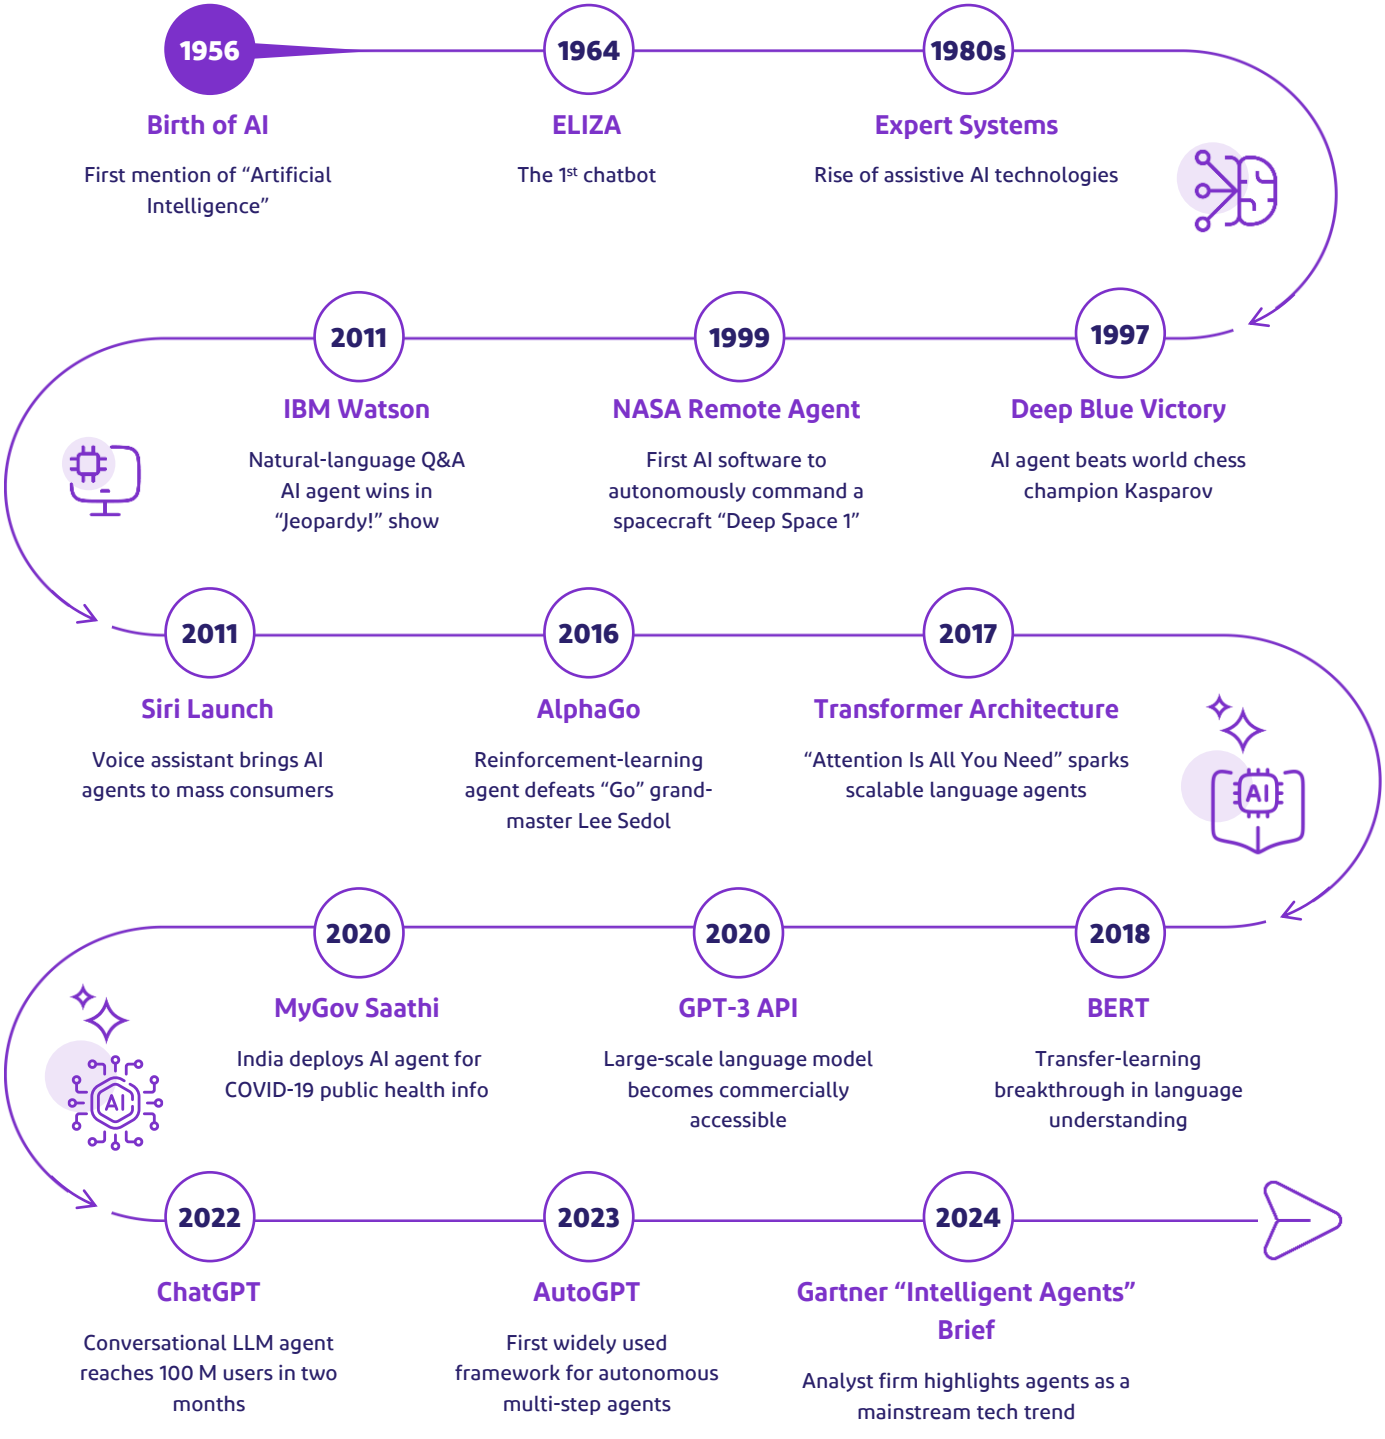


Figure 1: History of AI Agents as Government Partners (Timeline)

The evolution of AI agents follows a steady expansion of autonomy and intelligence. Foundational AI research at the 1956 Dartmouth workshop framed the idea of machines that could perceive and act. A decade later, ELIZA (1966) hinted at conversational interfaces, and the expert-system boom of the early 1980s put rule-based decision aids on government desktops, from tax-advice helpers to medical triage tools. These milestones proved that software could replicate narrow slices of expert reasoning and laid the groundwork for agents that assist rather than merely automate.

The 1990s and 2000s added strategic planning and real-time control to the agent toolkit. IBM's Deep Blue defeating Garry Kasparov in 1997 showcased machine strategy at super-human scale, while NASA's 1999 Remote Agent autonomously steered the Deep Space 1 probe, an early proof of agents operating in high-stakes, unsupervised settings. By the 2010s, breakthroughs in language and perception accelerated progress: IBM Watson's 2011 Jeopardy! win, the mass launch of voice assistants like Siri, AlphaGo's 2016 mastery of Go, and the 2017 Transformer architecture followed by BERT in 2018 collectively gave agents richer perception, reasoning, and learning capabilities.

Since 2020, large-scale generative models have pushed agents from research to mainstream governance. OpenAI's GPT-3 API (2020) offered on-demand language reasoning, while India's MyGov Saathi chatbot (2020) demonstrated crisis-response potential in the public sector. ChatGPT's viral 2022 launch, reaching 100 million users in two months, proved public appetite for conversational agents, and the open-source AutoGPT project in 2023 lowered barriers to building autonomous multi-step agents. Analyst houses now flag agentic AI as a top trend for 2024 and beyond, signaling that governments can begin road mapping from Basic conversational agents to Intermediate decision advisers and eventually to Advanced autonomous partners, balancing policy fit, cost, and oversight at each stage.

# 4. Significance of AI Agents

Artificial-intelligence agents are reshaping how software interacts with people, data, and public institutions. By perceiving information, reasoning over policy constraints, and acting on real-time feedback, modern agents can automate multi-step workflows and support decision-makers at scale. This section examines why agentic AI matters now, drawing on three vantage points: the global technology landscape, Saudi Arabia's national digital-government agenda, and examples of the public sector's need for adaptive AI agents.

## 4.1 Global Perspective

Benchmark evidence shows that agent capabilities are advancing quickly yet still fall short of dependable human-level performance. In 2024, researchers released VisualAgentBench, which evaluates embodied, GUI, and visual-design tasks. The top system, GPT-4o, completed only 36.2 percent of tasks, while most large language models hovered near 20 percent success, underscoring the difficulty of real-world agent operation. Another 2024 benchmark, RE-Bench, measures how well agents handle open-ended machine-learning R&D tasks. With a two-hour budget, leading agents scored roughly 4× higher than human experts, but humans overtook them at longer horizons, outperforming agents 2 to 1 after 32 hours demonstrating the challenge of maintaining effectiveness over the long term.

The AI Index 2025 adds further evidence of this trend. Large language model scores rose by nearly 50 points and more on challenging math and programming benchmarks, such as GPQA and SWE-bench. Yet the Index also shows that complex reasoning remains unsolved, with models struggling on structured logic tests such as PlanBench. Collectively, these results confirm rapid progress while highlighting the reliability gap that must be closed before agents can be trusted with high-stakes government duties.

Adoption pressures are rising in parallel with performance. Across 75 countries, parliamentary and congressional records logged a 21.3 percent increase in AI mentions during 2024, reflecting heightened attention to autonomous capabilities and governance. The United States alone issued 59 federal AI-related regulations, more than double the previous year, and 42 separate entities participated in rule-making. At the same time, governments are investing at scale: Canada announced a 2.4 billion-dollar AI infrastructure package, China launched a 47.5-billion-dollar semiconductor fund, France committed 117 billion dollars, India allocated 1.25 billion dollars, and Saudi Arabia's Project Transcendence allocates 100 billion dollars to national AI programs.

A recent example of governmental strategy is the U.S. White House's "America's AI Action Plan 2025". It outlines over 90 federal actions across three core pillars, accelerating innovation; building AI infrastructure; and leading international diplomacy and security.

These developments signal that agentic AI is no longer experimental. As models become cheaper to run, token inference costs for GPT-3.5-level performance fell from 20 dollars to 7 cents per million tokens in eighteen months, governments and enterprises have both the incentive and the budget to embed agents into daily operations. The global policy focus on transparency and safety, coupled with rapid performance gains, positions AI agents as a foundational capability for the next wave of digital transformation.

## 4.2 Local Perspective

Intelligent agents are positioned to be the next leap in Saudi Arabia's digital government. From 24-hour chat assistants to autonomous data-analysis tools, agents can shorten service times, strengthen decision quality, and free staff for higher-value work. Backed by Project Transcendence's one-hundred-billion-dollar AI investment and aligned with the National Digital Government Strategy (NDGS), the Kingdom now has the scale, governance principles, and roadmap to move from pilot bots to mission-critical adaptive agents across government entities.

### Citizens Satisfaction

AI agents elevate public services by offering personalized, real-time assistance across web, mobile, and call-centre channels. Conversational assistants can guide residents through license renewals, healthcare triage, and education enquiries, cutting wait times and ensuring clear, inclusive support in both Arabic and English.

### Enabling Businesses Sector

Agent-driven portals automate licensing, permitting, and regulatory queries so that entrepreneurs and small and medium-sized enterprises SMEs receive policy-aligned answers instantly. This reduces compliance costs, trims paperwork, and frees companies to focus on innovation and growth.

### Effective Governance

Policy-aware agents retrieve internal procedures, past decisions, and legal precedents on demand, helping officials reach faster, more consistent conclusions. Cross-agency agents can coordinate data flows, highlight conflicts, and recommend next steps, strengthening coherence across ministries.

## Effective Investment

Analytics agents mine national data assets to surface opportunities and risks in public-sector projects. By synthesizing market, labor, and infrastructure indicators, they support evidence-based capital allocation and more transparent evaluation of strategic initiatives. Those Agents enable decision-makers to effectively prioritize investments.

## Regulation And Compliance

Supervisory agents continuously scan filings, transactions, and open-source data for signs of non-compliance. They flag potential violations in near real time, generate audit trails, and propose corrective actions, enabling regulators to enforce standards swiftly and fairly.

## Expedited Transformation

Rolling out a tiered roadmap of basic, intermediate, and advanced AI agents accelerates government digital transformation. Integrated with existing portals and APIs, these agents break down information silos, scale automated support across departments, and align service delivery with NDGS targets.

# 4.3 Global Examples of AI Agents in the Government Sector

Intelligent agents are no longer confined to tech pilots; they are already carrying day-to-day workloads for public health hot-lines, legal-advice desks and citizen portals around the world. The four examples below show how widely the concept is spreading, the scale these agents can reach, and the diversity of missions they already support. Together they illustrate a practical roadmap for Saudi Arabia's government entities as they advance the use of AI agents as government partners.

## India

Flagship agent:
**MyGov Saathi** (Digital India)

24-hour multilingual chatbot on MyGov.in, WhatsApp and Facebook that pushes verified COVID-19 guidance, myth-busters and tailored advice to farmers, workers and seniors.

**Impact so far:** Handles ≈300 k queries per day / 20 k concurrent sessions and helped the "Corona Helpdesk" pass 30 million unique users during the pandemic.

## China

Flagship agent:
**Close Contact Detector** (National Health Commission)

AI-driven mobile service that cross-checks telecom mobility data with case registries so citizens can scan a QR code and see whether they were near a confirmed COVID-19 case.

**Impact so far:** Went live nationwide within 72 hours in Feb 2020 and became a core layer in China's data-centric outbreak containment stack.

## United Kingdom

Flagship agent:
**GOV.UK Chat** (Department for Science, Innovation & Technology)

GPT-4o agent that lets entrepreneurs ask plain-language questions (tax, licences, trade-marks) instead of navigating 700 000 web pages.

**Impact so far:** Second-phase trial opened to 15 000 users on 30 GOV.UK business pages; early tests saw ≈70 % of users rate answers helpful.

## Canada

Flagship agent:
**Botler AI** (DoJ-supported)

Conversational agent trained on 300 000+ Canadian & US legal documents that explains rights, drafts letters and routes citizens to resources-first focused on workplace-harassment cases.

**Impact so far:** Backed by Canada's Department of Justice in 2021 to expand nation-wide access to low-cost legal guidance, complementing human clinics and hotlines.

These examples show that well-designed agents can scale to tens of millions of users and still handle sensitive missions such as public health or legal guidance while maintaining high satisfaction rates. Each solution was built quickly by combining existing data with off-the-shelf AI tools, so cost and time barriers are low. Importantly, each deployment is overseen by dedicated regulators, (such as the UK's AI Safety Institute and Canada's Department of Justice). Strong governance keeps trust and accountability at the center of national deployments, offering a clear blueprint for Saudi Arabia's NDGS ambitions.

# 5. Methodological Framework and Success Factors for AI Agents' Deployment

AI agents deliver meaningful public-sector value when introduced through a structured approach. By perceiving data, making decisions, and acting in real time, agents can streamline services and strengthen policy execution. Moving from pilot projects to broad impact, however, requires clear frameworks. Successful programs align each agent's autonomy level (*basic, intermediate, or advanced*) with the use case, uphold transparency, and maintain strong human oversight. A practical framework helps entities scale responsibly, balancing efficiency gains with security and public trust. Such framework is presented below.

## Practical Framework and Critical Success Factors:

### 1. Technology Description

AI agents combine perception, reasoning, action, and learning within a single software entity. Capabilities range from Basic chatbots to Intermediate decision advisers and Advanced autonomous planners.

### 2. Adoption & Activation

Begin with low-risk, narrowly scoped agents and expand in stages. Match each deployment to policy fit, data maturity, and governance capacity, ensuring that humans remain in control of final decisions.

### 3. Potential Benefits

Faster citizen service, data-based policy support, and reallocation of staff to higher-value work. Agents can operate around the clock in multiple languages, improving accessibility.

### 4. Expected risks

Bias in decision outputs, data-security exposure, and over-dependence on automated actions. Entities must define reporting protocols and audit trails to keep agents accountable.

### 5. Critical Success Factors

High-quality data pipelines and robust model evaluation with ethical guidelines. Success also depends on a strong regulatory framework to guide adoption in the government sector, resilient AI infrastructure, and robust cyber-security measures.

### 6. Institutional Empowerment

Funding allocated to AI infrastructure, R&D, and secured data environments. entities should also invest in a skilled workforce through continuous technical training of experienced staff, ensuring human capacity evolves with technological progress.

This Methodological framework equips entities to introduce AI agents with confidence, protecting sensitive data and preserving human accountability. A disciplined rollout will let Saudi government entities unlock faster services, deeper insights, and stronger public trust while keeping oversight firmly in human hands.

# 6. Challenges for Government AI Agents Deployment

Deploying AI agents at national scale is a different order of complexity than running isolated chatbots. Saudi government entities already manage over 100 TB of information fed by 8,700+ datasets through the National Data Bank and Saudi Data Portal. Agents will need to read, write, and reason over this large amount of data without introducing new points of failure. They must also respect human-approval procedures and remain functional as APIs and legacy systems evolve. Reliability, supervision, and system compatibility therefore become as critical as model accuracy. This section highlights three challenges that require special attention (agent reliability, human-in-the-loop control, and platform interoperability) and suggests mitigation measures.

## 1 Agent Reliability and Error Cascades

### Key Concern

Autonomous agents often execute long chains of actions (planning, tool calls, and database writes) in a single workflow. A small error early in the chain can propagate and trigger large-scale failure. The AI Incident Database recorded 233 public AI incidents in 2024, a 56 percent rise over 2023, and 51 percent of organizations surveyed reported “unintended decision outcomes” when deploying AI. Such cascades erode trust and force costly efforts to fix mistakes.

### Mitigation

To contain cascade risk, entities should combine step-level validation with automatic rollback, so each action is verified before the next one begins. Observability is equally important: structured logs, confidence scores, and error codes must be captured so engineers and auditors can reconstruct every decision. Aligning these practices with ISO/IEC 42001 AI-management guidance and SDAIA’s national AI-ethics principles ensures reliability and auditability remain part of daily operations.

## 2. Autonomy Boundaries and Human Oversight

### Key Concern

Rising autonomy increases the likelihood an agent takes an unsupervised action in a high-stakes domain such as benefits approval or disbursement. Regulators are reacting: the United States issued 59 AI-related federal regulations in 2024, double the previous year, yet only 63 percent of executives report having formal oversight processes. Without clear hand-off points, entities risk policy misalignment, reputational damage, and diminished trust.

### Mitigation

Effective control starts with clearly defined human-in-the-loop checkpoints for every critical workflow, backed by RACI matrices that spell out who is Responsible, Accountable, Consulted, and Informed for each agent decision. Embedding these practices within the NIST AI Risk-Management Framework and OECD transparency principles helps ensure that oversight is systematic, auditable, and fully aligned with accountable digital government.

## 3. Interoperability and Dependency Risk

### Key Concern

Government AI agents must interact with 320 connected systems and 8,700 datasets already published through the National Data Bank and Saudi Data Portal, in addition to countless legacy applications and external APIs. Schema changes or decommissioned endpoints can strand agents or corrupt records, undermining the NDGS goals for seamless citizen-centric services.

### Mitigation

Sustained interoperability requires open interface standards such as OpenAPI, coupled with version contracts that alert engineers to breaking changes. Regression test suites should run automatically whenever an upstream API changes to catch silent failures. Leveraging shared infrastructure like SDAIA's Deem Cloud for central monitoring further ensures agents remain compatible as platforms evolve.

## 4. Data Privacy & PDPL Compliance

### Key Concern

AI agents aggregate and act on personal data across several entities. Without strict controls, they can exceed the original purpose, or expose records through logs, prompts, tool calls, or cross-border services. Under Saudi Personal Data Protection Law (PDPL), controllers must have a lawful basis, limit purpose and retention, honor data-subject rights, notify breaches per SDAIA guidance, and meet cross-border transfer conditions . Non-compliance risks penalties and loss of citizen trust.

### Mitigation

Adopt privacy-by-design for agents: enforce data minimization and role-based access; mask/redact PII in prompts and logs; localize processing where feasible; set automatic retention/deletion; and segregate agent contexts per case. Before deployment, run a Data Protection Impact Assessment (DPIA) and maintain a Record of Processing Activities (RoPA). For transfers, use SDAIA-approved SCCs or other safeguards and verify adequacy; avoid sending personal data to tools without a compliant transfer path.

## 5. Ethical AI and Bias Mitigation

### Key Concern

If AI agents inherit biased data or lack transparency, they may deliver outcomes that disadvantage certain groups or appear opaque to citizens. Without systematic auditing, errors or unfair patterns can persist undetected, eroding public confidence in government use of AI.

### Mitigation

Entities should adopt clear ethical AI frameworks with regular audits, fairness testing, and bias-detection methods. Independent review boards, explainability tools, and SDAIA's AI-ethics principles can ensure decisions remain transparent, accountable, and aligned with societal values.

# 7. Recommendations for Deploying AI Agents in government sector

Scaling AI agents from isolated pilots to mission-critical services requires more than writing code. Entities need shared infrastructure, Unified Governance Controls, and real-time visibility into how autonomous systems behave in production. The following recommendations offer a practical roadmap for government entities to embed reliable, interoperable, and accountable agents across government programs while advancing NDGS goals.

### Establish a National Agent Testbed and Simulation Lab

**Action:** Create a sandbox on SDAIA's Deem Cloud where entities can load synthetic or de-identified data, plug in candidate agents, and run stress tests that mimic real user traffic, edge cases, and adversarial attacks. Include scenarios like permit processing and disaster response, in addition to red-team tools to probe for bias or policy violations.

**Impact:** Provides a safe ground that surfaces reliability gaps before rollout, speeds cybersecurity accreditation, and builds a shared repository of best-practice evaluation scripts.

### Launch an Interoperability and Tool Registry Hub

**Action:** : Issue a central catalog, which could be managed by the DGA or SDAIA, that lists approved APIs, datasets, models, and software. Each entry should include version histories, performance benchmarks, and security approvals. Require agents to fetch tools through this registry so they automatically receive updates and deprecation notices.

**Impact:** Eliminates duplicate integration work, reduces breakage when APIs change, and ensures agents across 320 connected systems and 8,700 datasets can exchange information consistently and securely.

### Implement a Monitoring and Incident-Reporting Dashboard

**Action:** Deploy a cross-government platform that collects structured logs, confidence scores, and decision traces from every production agent. The dashboard should flag anomalies in real time, route incidents to designated response teams, and publish quarterly metrics on uptime, accuracy, and bias.

**Impact:** : Delivers continuous assurance that agents perform as intended, enables rapid remediation when issues arise, and strengthens public trust.

# 8. Conclusion

Artificial intelligence agents transform how digital government can operate. They route citizen requests to the right service in seconds, compress multi-agency decisions from weeks to hours, and keep hotlines open with always-on assistants that speak both Arabic and English. By reasoning across siloed data without moving it, agents promise a step-change in speed, accuracy, and inclusivity for public services.

This study has explained what makes an AI agent different from earlier automation efforts. It outlined Gartner's "perceive–decide–act–learn" cycle, introduced three capability tiers (Basic, Intermediate, and Advanced), and described how entities can map each tier to the policy fit, risk profile, and investment level that are right for them. A deployment framework linked these tiers to concrete benefits, from faster permit approvals to richer policy analysis.

The recommended roadmap is equally concrete. In includes practical recommendations such as A National Agent Testbed and Simulation Lab will expose reliability gaps in a safe environment. An Interoperability and Tool Registry Hub will give agents trusted access to the 320 connected systems and 8,700 datasets already available through the National Data Bank and the Saudi Data Portal. A government-wide monitoring and incident dashboard will provide real-time assurance on uptime, accuracy, and bias, strengthening public trust.

Global momentum makes timely action essential. Benchmarks such as VisualAgentBench show GPT-4o already solving 36 percent of complex tasks while most models hover near 20 percent, and RE-Bench reveals agents outperforming experts on short deadlines yet still trailing humans over longer horizons. Worldwide legislative mentions of AI rose 21 percent last year, the United States issued 59 federal AI regulations, and token inference prices for GPT-3.5-level performance fell from 20 dollars to seven cents per million tokens in just eighteen months. Inside the Kingdom, the 100-billion-dollar Project Transcendence and a seventh-place ranking in the UN E-Participation Index show that Saudi Arabia is ready to lead.

The DGA and other government entities like SDAIA are well placed to turn the testbed, registry, and dashboard into national standards, drawing on ISO 42001, the NIST AI Risk Management Framework, and local ethics guidelines. Successes with Absher, Tawakkalna, Najiz, and Deem Cloud offer a springboard for orchestrating agents at scale, setting KSA as a regional benchmark for accountable and efficient digital government.

The promise of AI agents goes beyond abstract algorithms and models. They are poised to deliver shorter wait times for a government services, faster answers for a parent seeking healthcare advice, and more transparent decisions overall for every resident query. Grounded in trusted national data and guided by strong oversight, AI agents can help Saudi Arabia deliver services that are timely, reliable, and aligned with Vision 2030, building public confidence in a truly citizen-centric digital future.

# 9. Definitions

| Term | Definition |
| --- | --- |
| AI Agent | An autonomous or semi-autonomous software system that can perceive information, make decisions, and take actions to achieve predefined goals. |
| Orchestration Agent | A type of AI agent designed to reason across multiple data sources or systems without moving the data itself. |
| Human-in-the-Loop (HITL) | A design approach where human oversight is built into the AI agent's workflow. This ensures that critical decisions, especially those involving policy, rights, or public trust are reviewed or approved by a human before execution. |
| Autonomy Level | A concept that describes the level of an agent's self-reliance in perception, decision-making, and execution. This level is determined by the agent's ability to handle risks and complexity of tasks, and the extent of the need for a human signature or approval. |
| Perceive–Decide–Act–Learn Cycle | A practical framework describing how AI agents function in real time: they perceive data from internal and external sources, decide on the best course of action, act through digital systems, and learn from outcomes to improve future performance. |
| Deem Cloud | A sovereign cloud platform developed by SDAIA that hosts secure government AI applications. |

# 10. Bibliography


1. Stanford Institute for Human-Centered Artificial Intelligence (HAI). (2025). Artificial Intelligence Index Report 2025. Stanford University. https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf
2. Gartner Research. (2025, June 25). When to use or not to use AI agents. Gartner.
3. Coshow, T. (2024, October 1). "Intelligent agents in AI really can work alone. Here's how." Gartner.
4. Weizenbaum, J. (1966). ELIZA – A computer program for the study of natural language communication between man and machine. Stanford University. https://web.stanford.edu/class/cs124/p36-weizenabaum.pdf
5. AutoGPT. (n.d.). AutoGPT: Autonomous AI agent platform. Retrieved from https://agpt.co
6. Saudi Data & AI Authority (SDAIA). (2024, September). State of AI in Saudi Arabia. Global AI Summit. https://globalaisummit.org/Documents/StateofAIinSaudiArabia.pdf
7. Digital India (MyGov). (n.d.). Saathi chatbot: AI-powered COVID-19 advisory and assistance. Retrieved from https://saathi.mygov.in
8. National Health Commission, People's Republic of China. (2020, February 10). Close Contact Detector: AI-driven mobile contact tracing service. Retrieved from https://en.nhc.gov.cn/2020-02/10/c_76428.htm
9. Department for Science, Innovation & Technology (UK). (2024). GOV.UK Chat: GPT-4-powered generative AI chatbot for business support. Retrieved from https://www.gov.uk/government/news/governments-experimental-ai-chatbot-to-help-people-set-up-small-businesses-and-find-support
10. Botler AI. (2021, March 3). Botler for Citizens: AI-powered legal assistance for workplace harassment (DoJ-supported). Department of Justice Canada. Retrieved from https://www.canada.ca/en/department-justice/news/2021/03/government-of-canada-supports-initiative-to-address-workplace-sexual-harassment-and-improve-access-to-justice-in-canada.html


**For continuous development purposes, please fill out the following survey**

Survey

SCAN HERE

**For more digital studies, please follow the link below**

Digital Studies

SCAN HERE

هيئة الحكومة الرقمية
Digital Government Authority

DGA.GOV.SA