




Agentic AI

Care Without Waiting

From pilot to governed autonomy: how agentic AI is reshaping the UAE health system toward services with no waiting and no friction.

Prepared by Eng. Loqman Omar Ali

Researcher in Agentic AI

logmanali@gmail.com

+971 581187347

June 2026

A strategic research document prepared to support decision-making in the UAE health sector. The analysis draws on official sources published by UAE government entities and on the latest research from leading global advisory firms (McKinsey, KPMG, Deloitte, PwC, EY) through June 2026.

The figures and estimates herein are indicative, referenced to their sources in the footnotes, and are not a substitute for a detailed feasibility study of each use case. This document does not constitute legal or medical advice.

© 2026 — All rights reserved to the author. No reproduction without attribution.

An in-depth research paper by **Eng. Loqman Omar Ali**, Researcher in Agentic AI, on turning agentic AI into a measurable operating program across the UAE health system — in full alignment with the "UAE Government 4.0" direction and the "Zero Bureaucracy" program.

Contents

Key Indicators	4
Executive Summary	5
01 The Inflection Point: Why Now	7
02 The National Frame: Government 4.0 & Zero Bureaucracy	10
03 Where the Value Lies in the Health System	14
04 The Operating Model: Governed Autonomy	18
05 Risk & Responsible Governance	22
06 The Roadmap & Decisions Required	25
07 The Playbook to Win the Zero Bureaucracy Award	29
Drug Registration — Ministry of Health & Prevention	31
Facility Licensing — Department of Health Abu Dhabi	31
Professional Licensing — Dubai Health Authority	32
Medical Appointment Booking — Emirates Health Services	33
08 The Practical Playbook: How to Start at Low Cost	35
Methodology & Sources	37

Key Indicators

- **50% of federal government services and operations** will shift to agentic AI within two years, under the "UAE Government 4.0" direction — making the UAE the first government in the world to adopt it at this scale.
- **80,000 federal employees** are enrolled in the largest training program in the government's history on agentic-AI tools, across five tracks from leadership to specialists.
- **Over 4,000 procedures eliminated and 12 million hours saved** — worth AED 1.12 billion — in the first phase of the "Zero Bureaucracy" program, cutting service time by more than 70%.
- **\$200–360 billion** in potential annual healthcare savings through AI; administrative burden alone consumes about 25% of health spending.¹
- **45.6% CAGR** for the global agentic-AI-in-health market, from \$0.54B to \$4.96B between 2024 and 2030.²
- **80% of healthcare leaders** expect agentic AI to deliver value, and 85% are raising their investment within 2–3 years.³
- **54% of organizations now deploy AI agents** in production (vs. 11% a year ago), yet only one in five has mature governance for them.⁴
- **41% reduction in clinical-documentation time** in live deployments (50% in leading systems) — time redirected straight back to patient care.⁵

¹McKinsey & Company, analysis of potential US healthcare savings. ²Grand View Research, agentic-AI-in-health market. ³Deloitte Center for Health Solutions, survey of 100 health organizations, Feb 2026. ⁴KPMG AI Pulse & Deloitte State of AI 2026. ⁵Mordor Intelligence; Deloitte (Nuance DAX).

Executive Summary

The UAE health system stands at a rare inflection point: the regulatory and institutional foundations are now ready to move from "piloting" AI to **operating it under governed autonomy** inside clinical workflows. The question is no longer "should we use AI?" but "how do we run it at scale without compromising patient safety, privacy, or trust?"

Agentic AI — also called agent-based AI — is the generation that moves beyond conversational assistants to systems that can **understand a goal, plan the steps, execute actions across multiple systems, and hand back to a human when needed**. In health, this means shifting from tools that "give a recommendation" to tools that "run a workflow": triaging cases, preparing approvals, coordinating appointments, following up after discharge, and linking data across records and operations.

This global wave intersects with an explicit national direction. In May 2026, the Cabinet approved the federal framework for the agentic-AI initiative, aiming to move **50% of government services and operations** to it within two years — alongside the "National Policy for Digital Health Services and AI in the Health Sector," built on three pillars: an AI-enabled national medical system, advanced digital-health infrastructure, and a qualified health workforce.¹

In parallel, the "Zero Bureaucracy" program set a decisive operating logic: eliminating over **4,000 procedures**, cutting service time by more than **70%**, and saving **12 million hours** worth AED 1.12 billion in its first phase alone.²

This is precisely where agentic AI becomes **the natural lever for the second phase** — moving past "digital bureaucracy" toward "government without friction and services without waiting."

The five advisory firms agree the shift has moved from experimentation to deployment: **54%** of organizations now run agents in production vs. 11% a year ago,³ **88%** of executives are raising budgets because of it,⁴ and **80%** of health leaders expect value from it.⁵ They also agree that **the barrier is not the technology but execution**: fragmented data, weak governance, and the skills gap.

This paper starts from a conviction that the highest value lies not in "medical advice," but in **high-frequency, high-friction health workflows**: emergencies, approvals, coordination, follow-up, and scheduling. These combine high impact with relatively low risk and precise measurability.

The paper proposes a **five-pillar national program** (governance, data, workflow, safety, skills), a smart matrix to prioritize use cases by impact and risk, an **autonomy ladder** linking the degree of automated delegation to the sensitivity of the clinical decision, and a

24-month operating roadmap aligned with the national target.

The bottom line

The UAE has what it takes to lead the region in health agentic AI. Success depends on **disciplined execution, not a proliferation of projects**: start fast with high-impact operational use cases, under strict governance and clear ROI measurement, then expand gradually toward deeper clinical uses as data, trust, and skills mature.

¹UAE Government Media Office, Cabinet meeting, 18 May 2026. ²UAE Government official portal, Zero Bureaucracy program. ³KPMG AI Quarterly Pulse. ⁴PwC AI Agent Survey 2025. ⁵Deloitte, Feb 2026.

Chapter One

The Inflection Point: Why Now

In twelve months, the world moved from generative AI to agentic AI capable of executing multi-step tasks. This chapter sizes the opportunity and the pace of adoption — and shows why 2026 is the inflection point, not the starting line.

The market is no longer nascent; it is growing faster than institutions can absorb it.


The agentic-AI-in-health market grows roughly ninefold in six years. More telling than the absolute number, institutional adoption has passed the pilot stage: organizations no longer ask whether it works, but how fast they can scale it under control.

45.6% CAGR of the global agentic-AI-in-health market (2024–2030)	\$505.6B Health-AI market size by 2033 (from \$36.7B in 2025)	\$2.6–4.4T Potential annual agentic-AI contribution to global output by 2030
--	---	--

Exhibit 1

The agentic-AI-in-health market expands nearly ninefold between 2024 and 2030.

Global market size, estimate, US\$ billion



~2026 and 2028 interpolated along the CAGR curve between the documented 2024 and 2030 points.

Source: Grand View Research, Agentic AI in Healthcare Market; author's analysis.

Self-Governing Institutions


From pilot to production: the curve flipped in a single year

The clearest evidence of market maturity is not its size but the behavior of organizations. In KPMG's quarterly surveys, the share of organizations actually deploying AI agents jumped from **11%** in early 2025 to **54%**, while those reaching systematic piloting passed **65%**. PwC confirms that **79%** of organizations use agents in some form, and **66%** of adopters record measurable productivity gains.

Exhibit 2

AI-agent deployment quintupled in a year — but mature governance fell behind.

Share of organizations deploying AI agents across the quarters of 2025, %



The critical gap

Despite the surge in deployment,
only one in five organizations

has a mature model for governing autonomous agents. This gap — not the technology — is what separates the "early adopters" from the "watchers."

Source: KPMG AI Quarterly Pulse Survey; PwC AI Agent Survey 2025; Deloitte State of AI in the Enterprise 2026.

Self-Governing Institutions

What this means for the UAE

The UAE is not entering an unknown market; it is joining a mature wave with documented ROI evidence. Its competitive edge is not in "trying" the technology — but in potentially being **the first to build mature governance at national scale**, which is precisely the missing link globally.

The National Frame: Government 4.0 & Zero Bureaucracy

What sets the UAE apart is that agentic AI is not an isolated initiative but a natural extension of an existing national frame: "UAE Government 4.0," the National Digital Health Policy, and the "Zero Bureaucracy" program. Agentic AI is the engine that turns these ambitions into real operations.

The UAE is on course to be the first government in the world to run half its services through agentic AI.

At the Cabinet meeting (May 2026), the federal framework defining the roles of ministries and entities in the agentic-AI initiative was approved, within a national transformation strategy aiming to move **50% of services and operations** to it within two years — a journey branded "UAE Government 4.0."

50%

of government services and operations move to agentic AI within two years

80,000

federal employees in the largest government training program on agentic-AI tools

4 cats.

citizen, resident, business, and public services in the first phase

This was accompanied by the "**National Policy for Digital Health Services and AI in the Health Sector**," which establishes a unified legal framework governing the development, licensing, accreditation, operation, and use of smart health systems. H.H. Sheikh Mansour bin Zayed leads and oversees this journey.

Exhibit 3

The National Digital Health Policy rests on three integrated pillars — none of which works in isolation.

The three approved pillars, 2026

01

An AI-enabled national medical system

An integrated health system supported by AI models across service pathways, from emergencies to follow-up.

02

Advanced digital-health infrastructure

A trusted, traceable, interoperable data and operations foundation — the precondition without which no agent succeeds.

03

A health workforce skilled in AI

Physicians, nurses, and administrators who understand agentic AI's capabilities and limits — the shift is as cultural as it is technical.

Source: UAE Government Media Office, Cabinet meeting, 18 May 2026; author's analysis.

Self-Governing Institutions

"Zero Bureaucracy": the logic that turns agentic AI into a reform tool

The "Zero Bureaucracy" program proved that re-engineering procedures yields huge, measurable value. In its second phase, the focus shifted to "digital bureaucracy" and embedding AI, with a stated goal: **"government without friction, services without waiting."** Agentic AI is the tool that delivers this goal in the health sector specifically — where waiting is not just an administrative nuisance but a clinical risk.


Exhibit 4

Phase one of "Zero Bureaucracy" saved 12 million hours worth AED 1.12 billion — before smart agents even arrived.

Phase-one results, across 30 entities and 690 work teams



Service completion time, before vs. after



Source: UAE Government official portal; Media Office, launch of Zero Bureaucracy phase two, June 2025.

Self-Governing Institutions

The core idea

"Zero Bureaucracy" removed redundant procedures *manually*. Agentic AI completes the task *automatically*: it does not just delete the step but **executes it on the human's behalf** — turning "simplifying the procedure" into truly "zeroing it out" of the patient's and employee's experience.

A four-layer system: a multi-level national design

The structure of the UAE health sector enables a multi-layered national model rather than a single standalone pilot. Each layer has a distinct role that reinforces integration, not duplication.

Exhibit 5

Four regulatory layers, four complementary roles — federal enablement meets advanced governance and scalable delivery.

Map of UAE health entities and each one's role in the transformation

Ministry of Health & Prevention	Federal enablement & regulation The first federal center of excellence for AI in health; and the PaCE smart-operations center for emergency-flow management and prediction — winner of 10+ global awards.
Department of Health — Abu Dhabi	The responsible-governance model A "Responsible AI" standard mandating transparency, explainability, and human intervention; and the Global Healthcare AI Academy (with MBZUAI and Core42).
Dubai Health Authority	Customer-journey integration Embedding agentic AI into the patient, hospital, and clinical/administrative networks — across planning, operations, and service.
Emirates Health Services	Scalable operational delivery Live smart solutions (five showcased at Arab Health 2025) in readmission prediction, process optimization, and resource management.

Source: official MoHAP, DoH, DHA, and EHS websites; author's analysis.

Self-Governing Institutions

Chapter Three

Where the Value Lies in the Health System

The highest value of agentic AI is not in "general medical advice," but in high-frequency, high-friction health workflows. This chapter offers a smart matrix to prioritize use cases, identifies the value pools, and maps the six service pathways.


Start where high impact meets low risk — not where clinical sensitivity concentrates.

The matrix below plots health use cases on two axes: operational impact (vertical) and clinical sensitivity/risk (horizontal), with bubble size reflecting process frequency. The result is a national priority list in three waves.

Exhibit 6

Seven operational use cases fall in the "start zone": high impact, low risk, high frequency.

Priority matrix: operational impact vs. clinical sensitivity; bubble size = process frequency



Source: author's work, drawing on McKinsey, Deloitte, KPMG, and EY on high-value health use cases.

Self-Governing Institutions

Value pools: administrative burden is the big prize

McKinsey estimates that broad AI adoption could unlock **\$200–360 billion** in annual healthcare savings. Paradoxically, most of it comes not from diagnosis but from **automating administrative burden**, which alone consumes about a quarter of health spending.

Exhibit 7

Five value pools — the largest are administrative automation and revenue cycle, not clinical decisions.

Potential healthcare savings pools, ranked by estimated relative impact



Executive read

Hospitals could cut costs by **4–11%** and physician groups by **3–8%** — mostly via process improvement and administrative automation, not direct clinical intervention. This confirms the "start with operations" logic.

Source: McKinsey & Company, US healthcare savings estimates; author's analysis. Bar display is relative/indicative.

Self-Governing Institutions

The six service pathways: where each delivers the most impact

Value spreads across six main service pathways, each with varying impact across five dimensions. The heatmap below shows emergency and approvals leading on breadth of impact, while assisted documentation stands out specifically in staff productivity.

Exhibit 8

Emergency and approvals deliver broad impact across all dimensions; assisted documentation excels in staff productivity.

Impact intensity per service pathway across five dimensions (darker = higher impact)

	Wait time	Cost	Customer experience	Staff productivity	Patient safety
Emergency & flow	High	Med	High	Med	High
Approvals & claims	High	High	Med	High	Low
Coordination & referrals	Med	Med	High	Med	Med
Post-discharge follow-up	Low	High	High	Med	High
Scheduling & resources	High	Med	Med	High	Low
Assisted clinical documentation	Low	Med	Med	High	High

Assisted documentation cuts clinical-documentation time by 41% in live deployments (50% in leading systems), freeing staff time for direct care.
Source: author's work, drawing on the five advisory firms' reports and documented deployments.

Self-Governing Institutions

Chapter Four

The Operating Model: Governed Autonomy

The solution is not a tool, but a five-pillar national program:

governance, data, workflow, safety, and skills. This chapter presents the operating architecture, the logic of "zeroing out" manual steps, and the autonomy ladder that links automated delegation to decision sensitivity.

Five pillars, none of which works alone — governance first, then data, then deployment.

The five advisory firms agree that projects fail because of fragmented data, weak governance, and the skills gap — not the models. So the national program starts with the foundation, not the shiny application.

Exhibit 9

The national program for agentic AI in health: five priority-sequenced pillars.

The operating architecture, from foundation to scale

01	Governance before scale A unified federal-local framework defining: who approves, who reviews, when a human intervenes, and the limits of automated decisions — with mandatory audit logs and stop authority.
02	Unifying the data layer & interoperability Shared standards, clinical and operational data dictionaries, and linking records to metrics — no agent succeeds atop untrusted or untraceable data.
03	Re-engineering workflows (not just digitizing them) Selecting high-impact, low-risk use cases and zeroing out redundant steps before automating them — extending the "Zero Bureaucracy" logic.
04	Safety & "human in the loop" Pre-deployment testing, drift monitoring, quality and safety metrics, and human intervention in every clinically impactful use.
05	Skills & professional accreditation A national program covering physicians, nurses, administrators, and tech teams — aligned with the 80,000-employee federal training program.

Source: author's work, aligning KPMG (the real barriers), DoH Abu Dhabi (the Responsible AI standard), and the National Digital Health Policy.

Self-Governing Institutions

"Zero Bureaucracy" in practice: from a manual journey to a touchless path

Take one example: **pre-authorization for a medical procedure**. In the traditional model it passes through a chain of manual steps between patient, physician, and insurer — taking days and generating friction and errors. The agent does not merely "simplify" the step — it **executes it**: reads the record, assembles documents, checks the criteria, and issues the decision, with human escalation only on exceptions.

Exhibit 10

Re-engineering pre-authorization cuts 12 manual steps to 3, turning days into minutes.

Illustrative model of the pre-authorization journey, before and after the agent



Illustrative model combining documented deployment patterns in approvals and claims across health-tech vendors; ratios are indicative.

Source: author's work, drawing on McKinsey (the race to a touchless revenue cycle) and PwC.

Self-Governing Institutions

The generalizable principle

Each of the six pathways has its own "version" of this journey. The goal is not to add an intelligence layer atop a complex system, but to **remove the complexity, then execute it automatically** — the essence of "services without waiting."

The autonomy ladder: automated delegation is inversely proportional to clinical sensitivity

The elegant answer to the "trust vs. speed" dilemma is not to grant every case the same degree of autonomy. A five-level ladder links the delegation granted to the agent to the sensitivity of the decision: near-autonomous operation for low-risk tasks, full human oversight for clinical decisions.

Exhibit 11

Five levels of autonomy — the higher the clinical sensitivity, the tighter the human oversight.

The autonomy ladder and human-escalation thresholds

L4	Autonomous execution with post-hoc audit — scheduling, reminders, queue management. The agent executes and logs; review is by sampling. Suited to purely administrative tasks.
L3	Execute with notification and override — routine approvals, refills. The agent executes and notifies the official, who can override within a time window.
L2	Proposal requiring human approval — referral triage, follow-up plans. The agent prepares the full decision; a human approves before execution.
L1	Advisory support only — diagnostic decision support, risk alerts. The agent suggests; the decision and responsibility rest fully with the human.
L0	Fully human — high-risk clinical decisions. No automated delegation; AI is outside the decision loop.

The governance rule A use case moves up the ladder only after clearing performance and safety thresholds and obtaining regulator approval — never by default. This turns "autonomy" from a risk into an earned privilege.
--

Source: author's work, aligning the Responsible AI standard — Department of Health Abu Dhabi, and EY's role-redistribution framework.

Self-Governing Institutions

Chapter Five

Risk & Responsible Governance

Agentic AI without strict governance is just a technical layer atop inconsistent systems — multiplying risk instead of unlocking value. This chapter maps the risk register and its mitigations, and the governance model aligned with Abu Dhabi's Responsible AI standard.

The risks are known and managed — mitigation starts at design, not after launch.

Every material risk has a specific mitigation built into the system from the start. The table below ranks the seven leading risks by likelihood and impact, pairing each with its practical mitigation.

Exhibit 12
Seven material risks — data fragmentation and cybersecurity rank highest, each with mitigation built into the design.

Risk register: likelihood × impact, with mitigation (darker = higher)

	Likelihood	Impact	Built-in mitigation
Data fragmentation & weak interoperability	High	High	Unified data layer, interoperability standards, shared clinical dictionaries
Cybersecurity & data leakage	Med	High	Sovereign architecture, encryption, local processing of sensitive data, no production data in training
Bias & weak clinical explainability	Med	High	Fairness testing, decision-logic explanation, periodic model review
Undefined accountability on error	Med	High	Clear accountability framework, mandatory audit logs, stop authority
Vendor over-dependence	High	Med	Strict partner-compliance requirements, switchability, entity data ownership
Blind automation & over-reliance	Med	Med	"Human in the loop," escalation thresholds, a graduated autonomy ladder
Change resistance & skills gap	High	Med	National skills program, involving staff in design from the start

Source: author's work, aligning KPMG (the future of AI governance) and the DoH Abu Dhabi standard.

Self-Governing Institutions

The governance model: Abu Dhabi's Responsible standard as a national foundation

Abu Dhabi's "Responsible AI" standard offers a ready basis for national rollout. It imposes on every licensed facility a set of obligations that protect the patient without stifling innovation.

Requirements of the Responsible standard

Data-source documentation: full traceability of every decision's origin.

Decision-logic explanation: explainability for clinician and patient.

Patient notification: disclosure when AI is used.

Right to opt out: available where appropriate.

Human intervention & escalation: clear procedures on exceptions.

Risk & security management: continuous compliance before and after deployment.

The governing principle: aligning federal and local frameworks

The biggest governance risk is not its absence but its **divergence** across entities. The next phase therefore requires a unified framework that prevents compliance requirements from differing facility to facility — preserving the pace of scaling without fragmenting accountability. The Department of Health has already begun coordinating with international organizations to draft principles for AI governance in healthcare.

"The challenge is no longer: should we use AI? But: how do we run it at scale without harming safety, privacy, or trust?"

— A shared conclusion from KPMG and Deloitte surveys on AI governance, 2025–2026

The Roadmap & Decisions Required

The gap is no longer in vision or technology, but in **the tempo of execution**. This chapter translates the program into a four-quarter roadmap aligned with the national target — 50% of services through agentic AI within two years — and sets out the national scorecard and the six decisions leadership must make now.


Four quarters, not a linear sequence: the foundation is built early, and the waves overlap rather than wait.

Stalled transformations make the mistake of sequencing the work: governance, then data, then deployment. A successful roadmap **overlaps** the workstreams — governance and data start together in Q1, the operational wave launches before the infrastructure is complete, and the skills program runs across the full two years.

Exhibit 13

A 24-month roadmap: six overlapping workstreams peaking at the national 50% target.

Distribution of workstreams across four time quarters, aligned with the federal Government 4.0 framework



Piloting begins at selected sites across federal and local entities, then scales after consolidating lessons learned. Each wave passes through the scorecard before moving to the next — no scaling without evidence of impact.

Source: author's work, aligned with the federal agentic-AI framework (Cabinet, May 2026) and McKinsey and Deloitte transformation roadmaps.

Self-Governing Institutions

The national scorecard: no scaling without evidence of impact

Every workstream in the roadmap is tied to a quantitative metric measured before and after scaling. This scorecard turns agentic AI from a "tech project" into an **institutional reform tool** with proven ROI — embodying the "Zero Bureaucracy" principle: from multiple manual steps to a near-autonomous path.

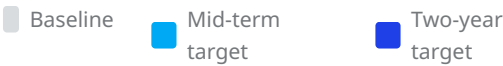


Exhibit 14

The national performance card: six metrics that govern the move from one wave to the next.

Baseline, 12-month target, and 24-month target — each metric with a designated business owner

National metric	Baseline	Target 12 mo	Target 24 mo	Dimension
Emergency wait time	Current ref.	-35%	-70%	Patient experience
Pre-authorization cycle time	Days	Hours	Minutes	Operational efficiency
Readmission rate (30 days)	Current ref.	-15%	-30%	Clinical quality
Administrative hours saved annually	—	Millions	Tens of millions	Financial return
Governed-autonomy coverage	<5%	20%	50%	Transformation maturity
Patient & staff satisfaction	Current ref.	+10 pts	+22 pts	Trust & adoption



The +22 pts satisfaction value derives from documented results for clinical-documentation assistants (Nuance DAX) that cut documentation time by about 50%. Figures are indicative targets, calibrated to each entity's actual baseline.
Source: author's work, aligning MoHAP (PaCE center) metrics, "Zero Bureaucracy," and Deloitte health surveys 2026.

Self-Governing Institutions

The six decisions leadership must make now

Moving from "scattered piloting" to "disciplined transformation" requires six interlinked executive decisions. Each is a necessary condition; together they form the institutional contract for a national program able to reach the two-year target.

01	Adopt agentic AI as a national health priority Embed it within government digital programs and Government 4.0, with an explicit executive mandate and multi-year budget — not as an isolated pilot.
02	Task a federal body with coordination and unified governance Setting shared compliance standards with local entities, built on Abu Dhabi's Responsible standard, to prevent diverging requirements facility to facility.
03	Approve a national list of priority use cases Five to seven cases starting from the operational "start zone," ranked by impact and risk, each with a designated business owner.
04	Fund the data layer and interoperability before scaling Unifying standards, clinical and operational dictionaries, and linking records to metrics — agentic AI does not succeed atop fragmented, untraceable data.
05	Mandate "human-in-the-loop" design as a binding standard In every application touching the patient or the clinical decision, with pre-deployment testing, drift monitoring, stop authority, and audit logs.
06	Launch the skills and professional-accreditation academy Qualifying leaders, physicians, nurses, administrators, and tech teams — the shift is operational and cultural as much as technical, and skills are the biggest barrier to scaling.

The bottom line

The UAE has the regulatory and institutional foundations to lead the region in health agentic AI. **Success is measured not by the number of projects, but by the discipline of execution:** start fast with high-impact operational cases, under strict governance and clear ROI measurement, then expand gradually toward the more complex as data, trust, and skills mature. This is how AI turns from a technical slogan into **care without waiting**.

The Playbook to Win the Zero Bureaucracy Award

Why these services specifically? Selection criteria per the Zero Bureaucracy rubric.

The services were not chosen at random, but per the **official evaluation axes** of the "Zero Government Bureaucracy" program (Cabinet Office) and its stated targets: **cutting at least 50% of processing time**, zeroing out unnecessary requirements, prioritizing the procedures **that most affect people's lives and businesses**.

The program's official evaluation axes

Impact How much the procedure affects people's lives and businesses.	Engaging entities & individuals Involving customers and employees in the redesign.
Compliance with targets Achieving $\geq 50\%$ time cuts and zeroing out requirements.	Empowering employees Freeing staff from repetitive manual work.

Priority matrix — why each service was selected

Selected service	Impact	Duration (reduction opportunity)	Complexity	Usage intensity
Drug registration — MoHAP	●●● National	●●● 30–90 days	●●●	●●
Facility licensing — DoH Abu Dhabi	●●● Sector growth	●●● ~45 working days	●●●	●●
Professional licensing — DHA Dubai	●●● Workforce	●● 2–3 months	●●	●●● Thousands/yr
Appointment booking — EHS	●●● Every patient	●● Daily, recurring	●●● Complex rules	●●● Largest daily

Source: UAE Government official portal (u.ae), Ministry of Finance, and the Cabinet Office — axes and targets of the Zero Government Bureaucracy program; author's framing. The scale (●) is relative/indicative.

Ministry of Health & Prevention

Zero Bureaucracy · Agentic AI

Longest service: Pharmaceutical / drug product registration

Why this service? Highest **national impact** (drug availability), and the longest **duration** (30–90 days) — i.e. the biggest opportunity to hit the $\geq 50\%$ reduction target — with **high complexity** in dossiers and requirements; a perfect match to the rubric's priorities.

Current state: four sequential stages — technical dossier review (CTD) → lab analysis → registration-committee review → pricing approval; with query correspondence that stops the clock and extends the timeline.

Agent activation phases — gradual

Phase 1 · Smart triage & completeness (employee copilot)

The agent receives the dossier, automatically checks section and document completeness against Federal Law 8/2019, and instantly generates the gaps-and-queries list instead of weeks of manual review — the employee approves.

Phase 2 · Augmented technical review (human in the loop)

The agent extracts critical data (composition, stability, test results), compares it with the knowledge base and registered products, and escalates risks and alerts to the human reviewer who decides.

Phase 3 · Committee & pricing (autonomous for simple cases)

The agent prepares the committee file, proposes internationally benchmarked reference pricing, and automates certificate issuance and tracking for low-risk products.

Toward winning: from 30–90 days to a handful of days (a reduction beyond the 50% target), with an explainable decision and real-time tracking.

Department of Health — Abu Dhabi

Zero Bureaucracy · Agentic AI

Longest service: New healthcare facility license

Why this service? Direct impact on **sector growth and investment**, a long duration (preliminary approval 20 + permanent license 25 working days), and **high complexity** across multiple entities — an ideal opportunity to meet the time target.

Current state: via TAMM — economic approvals → design & drawings → inspection → permanent license; the bottleneck: document exchange across multiple entities and manual review of drawings.

Agent activation phases — gradual

Phase 1 · Instant preliminary approval

The agent automatically checks requirements and economic approvals, guides the investor with a precise requirements list, and issues preliminary approval in hours instead of 20 working days.

Phase 2 · Smart drawing review & inspection

The agent visually reviews drawings against the Health Facility Guidelines (HFG) and flags violations, auto-schedules inspection and prepares its report — with a human inspector for the final decision.

Phase 3 · Touchless permanent licensing

The agent links inspection results to automatic permanent-license issuance via TAMM, and tracks status in real time for the customer — aligned with the Department's "Responsible AI" standard.

Toward winning: from ~45 working days to a few days, with explainable, responsible governance — directly accelerating sector growth.

Dubai Health Authority

Zero Bureaucracy · Agentic AI

Longest service: Health professional license

Why this service? Highest **usage intensity** (thousands of professionals per year) and direct impact on the **health workforce**, with a long duration (2–3 months) caused by a clear bottleneck — the highest return per dirham invested.

Current state: primary-source verification (DataFlow / PSV) 6–8 weeks → assessment & approval 1–2 weeks → exam → activation; the bottleneck: manual verification of documents from their sources.

Agent activation phases — gradual

Phase 1 · Parallel upload & verification

The agent extracts certificate data and starts verification (DataFlow) automatically and in parallel across sources, tracks responses and reminds the sources — compressing 6–8 weeks into days.

Phase 2 · Equivalency assessment & exam scheduling

The agent assesses qualifications against specialty criteria, determines the need for an exam and schedules it automatically, and escalates borderline cases to the human assessor.

Phase 3 · Instant activation via "Sheryan"

On acceptance and facility linkage, the agent activates the license and issues the card digitally on the spot — a natural extension of a platform that already uses AI.

Toward winning: from 2–3 months to a week, in the most frequently used service — the highest impact per dirham.

The common thread across the four pathways

One structure recurs in every service:

automated completeness checking

first, then

AI-augmented review with a human in the loop

, then

autonomous handling of low-risk cases

. This same progression is what delivers the $\geq 50\%$ time-reduction target without compromising governance — turning the employee from a manual processor into a supervisor of exceptions.

Emirates Health Services · The Pivotal Case

Medical appointment booking: from manual rule-chaos to a correct booking on the first try.

The governing thought: the most intensive and complex service in the organization is the highest-return candidate for automation. A single agent — absorbing all the rules and integrating with "Wareed" — zeroes out booking errors and cuts call time beyond the program target ($\geq 50\%$), gradually proving ROI before every expansion.

Why this service? (Zero Bureaucracy rubric): highest **usage intensity** (the largest daily) and **highest rule complexity**, with impact touching **every patient** — a direct priority for zeroing out errors and cutting time.

Exhibit 15

Three interwoven dimensions intersect in every call — and from their intersection the error arises.

Anatomy of medical-appointment-booking complexity





Exhibit 16

The magic wand: one rules engine that replaces the fragile manual juggling.



A gradual rollout map — proving ROI at each stage before scaling.

The shift requires no single leap, but four stages — each starting as human assistance and ending in broader automation, with the human staying in the loop until each pathway matures. This gradual approach is what reassures leadership and meets the time-reduction target without compromising governance.

Exhibit 17

From "employee copilot" to "touchless" — four stages, each with tangible proof.

Stage	What happens	Technical integration	Human role	Evidence / metric
0 Prove the value	One low-risk specialty (general medicine/follow-up): the agent suggests the right slot and the employee decides.	Read availability from "Wareed" (read-only).	Employee copilot.	–booking errors & –call time within 8–12 weeks.
1 Absorb the complexity	Add the complex-pathway rules (dentistry, obstetrics, follow-up, referral) and doctors' hours and availability.	Full integration with "Wareed" + CRM alignment to unify the customer record.	Human in the loop for unusual cases.	Accurate coverage of all pathways.
2 Guided self-booking	The customer books directly while the agent applies the rules behind the scenes and writes the booking instantly.	Across the EHS app , website, voice channel, and call center, writing into "Wareed."	Escalate complex cases only.	Reduced call-center load.
3 "Touchless" operation	Full booking management across all channels and specialties, with rescheduling and reminders.	Sync "Wareed" & CRM across all channels.	Supervisor of exceptions.	Correct booking on the first try.

Durations and metrics are indicative and calibrated to the organization's data.

THE DECISION REQUIRED FROM THE CHIEF AI OFFICER

Approve **Stage 0** only for now: a ring-fenced pilot on **one specialty** for **8–12 weeks**, where the agent is an **employee copilot** reading from "Wareed" with no write access — near-zero risk. The criterion: a tangible drop in booking errors and call time opens the door to the next stage. **A small decision, proven ROI, risk-free expansion.**

The outcome: a correct booking on the first try, a radical cut in errors, and call-time reduction beyond 50% (the program target) — with both customer and employee satisfaction.

Source: official entity websites (ehs.gov.ae, mohap.gov.ae, doh.gov.ae / TAMM, dha.gov.ae / Sheryan); the appointment-booking service and the EHS app integrated with the "Wareed" health information system (ehs.gov.ae/about-us/smart-applications-of-ehs); and the Zero Bureaucracy program axes (u.ae); author's framing. Durations are indicative and subject to change.

An executive framework · a practical guide for the decision-maker

The practical playbook: the methodology and the low cost to start.

McKinsey research shows value comes not from scattered pilots but from **"prove it, then scale it"**: start where the return is immediate, delegate a full task to the agent (documentation, patient communication, scheduling), and only scale after validation. The big dilemma is not technical: **75%** of US hospitals run at least one AI application, yet **only 8%** have moved beyond a single department or small pilot. Discipline in scope — not budget size — is the difference.

The methodology in six stages

01	02	03	04	05	06
Strategy & alignment	Use-case selection	Proof of concept (pilot)	Architecture & integration	Controlled deployment	Governance & scale

Financial estimates to start — deliberately low

Value can be proven on a modest budget before any large institutional commitment. Figures are indicative, in US dollars.

Track	Scope & output	Duration	Estimated cost
Basic agent (RAG)	Chat/retrieval assistant for one specific task	4–8 weeks	\$8–25k (run \$0.5–2k/mo)
Proof of concept (PoC)	One case, isolated environment, model evaluation — a "go / no-go" decision	6–12 weeks	\$50–100k
Controlled pilot	Real data, a small team, one or two system integrations, proven ROI	3–6 months	\$50–200k
Year one for a health facility	One to two cases through to live operation with governance	12–18 months	\$75–500k

Source: author's compilation from McKinsey (Pilot-to-Platform) and 2026 health-sector implementation cost guides (Azilen, Integrio, Interexy, Riseup). Costs fall 30–50% when using a specialist partner instead of building in-house.

Cost-control guardrails — where projects stumble

- **Clean the data first:** fragmented data can consume up to 40% of the budget; allocate \$5–50k for de-identification and governance before any training.
- **Govern "shadow AI":** ungoverned tools are present in 40% of hospitals and add ~\$670k to the cost of any breach — the governance layer is an investment, not a luxury.
- **Add a contingency ≥20%** and favor a specialist partner (saving 30–50%) over building in-house at the start.

The 90-day map: from decision to proven ROI.

The six essentials — before you start

01

An executive sponsor

Leadership that owns the decision and resources and removes blockers.

02

One use case

From the "start zone": high impact, low risk, high frequency.

03

Linkable data

A clean source and a unified dictionary; preparing it can swallow 40% of the budget.

04

A small team

Business + tech + compliance, with clear responsibilities.

05

Minimum governance

Limits of automated decisions, intervention points, stop authority, and an audit log.

06

One success metric

A measurable number on which the scaling decision is made.

The three sprints — the 90-day map

Days 1–30

Foundation & selection

- Pick one case: documentation, patient communication, scheduling, or the revenue cycle.
- Assign the sponsor, business owner, and governance lead.
- Measure the baseline: time, cost, errors, satisfaction.
- Approve a proof-of-concept budget of \$50–100k.

Days 31–60

Build & integrate

- Connect the data source, unify its dictionary, and clean it first.
- Design the "touchless" workflow with a human in the loop.
- Test in an isolated environment with audit logs.
- Train the team on oversight and escalation.

Days 61–90

Limited launch & measurement

- Launch a live, narrow-scope pilot with one team.
- Measure weekly: 40–60% time reduction, adoption after 4 weeks.
- Correct drift and expand gradually.
- Make the scaling decision based on ROI evidence.

The golden rule & ROI evidence

Start with one case, prove the ROI, then scale. The evidence is strong: **64%** of implementing organizations report positive ROI, and one case study documented **\$40 million** in contribution margin from AI operating-room scheduling alone. Start small and stay disciplined in scope — and from the accumulation of proven cases, the organization reaches the national 50% target.

Source: McKinsey (Prove it, then scale it; the operating-room scheduling case), and 2025–2026 health-AI cost and performance surveys.

Methodology

This paper combines **analysis of official UAE sources** with **the latest research from the five global advisory firms** through June 2026. A three-layer analytical framework was used: (1) mapping the governing national directions — the federal agentic-AI framework, Government 4.0, the "Zero Bureaucracy" program, and the National Digital Health Policy; (2) deriving value pools and use cases from the quantitative surveys of McKinsey, KPMG, Deloitte, PwC, and EY; (3) aligning the two into a measurable operating model via the impact/risk matrix, the autonomy ladder, and the national performance card.

The target figures in Chapter Six are **indicative**, calibrated to each entity's actual baseline before adoption. The paper contains no quotes attributed to specific individuals; conclusions are attributed to their organizations and published surveys.

Primary sources

Official UAE entities

- Cabinet Office — meeting of 18 May 2026 (the federal agentic-AI framework and the National Digital Health Policy).
- Ministry of Health & Prevention — first federal AI center of excellence; the PaCE smart-operations center.
- Department of Health – Abu Dhabi — the Responsible AI standard, 2025.
- Emirates Health Services — the five smart solutions, Arab Health 2025.
- The "Zero Government Bureaucracy" program — phase-one and phase-two results.

Global advisory firms

- McKinsey & Company — agentic-AI research and healthcare value pools.
- KPMG — the quarterly "AI Agent Pulse" surveys and the future of AI governance.
- Deloitte — the agentic-AI healthcare survey, 2026.
- PwC — agentic-AI outlook and economic impact.
- EY — redistributing roles between humans and technology in care.

A methodological note

The statistics here represent the latest publicly available data at the time of preparation and may change as surveys evolve. Their purpose is to build a coherent strategic argument, not to provide numerical guarantees. Executive decisions remain subject to the judgment of the competent entities and their actual data.

AI

Self-Governing Institutions Series

Strategic Research Paper · June 2026



LOGMAN OMER ALI

مستشار التخطيط الاستراتيجي
وباحث في الذكاء الاصطناعي المساعد

Logman Omer Ali



Scan to connect on LinkedIn

Contact

logmanali@gmail.com

+971 581187347

linkedin.com/in/logmanali