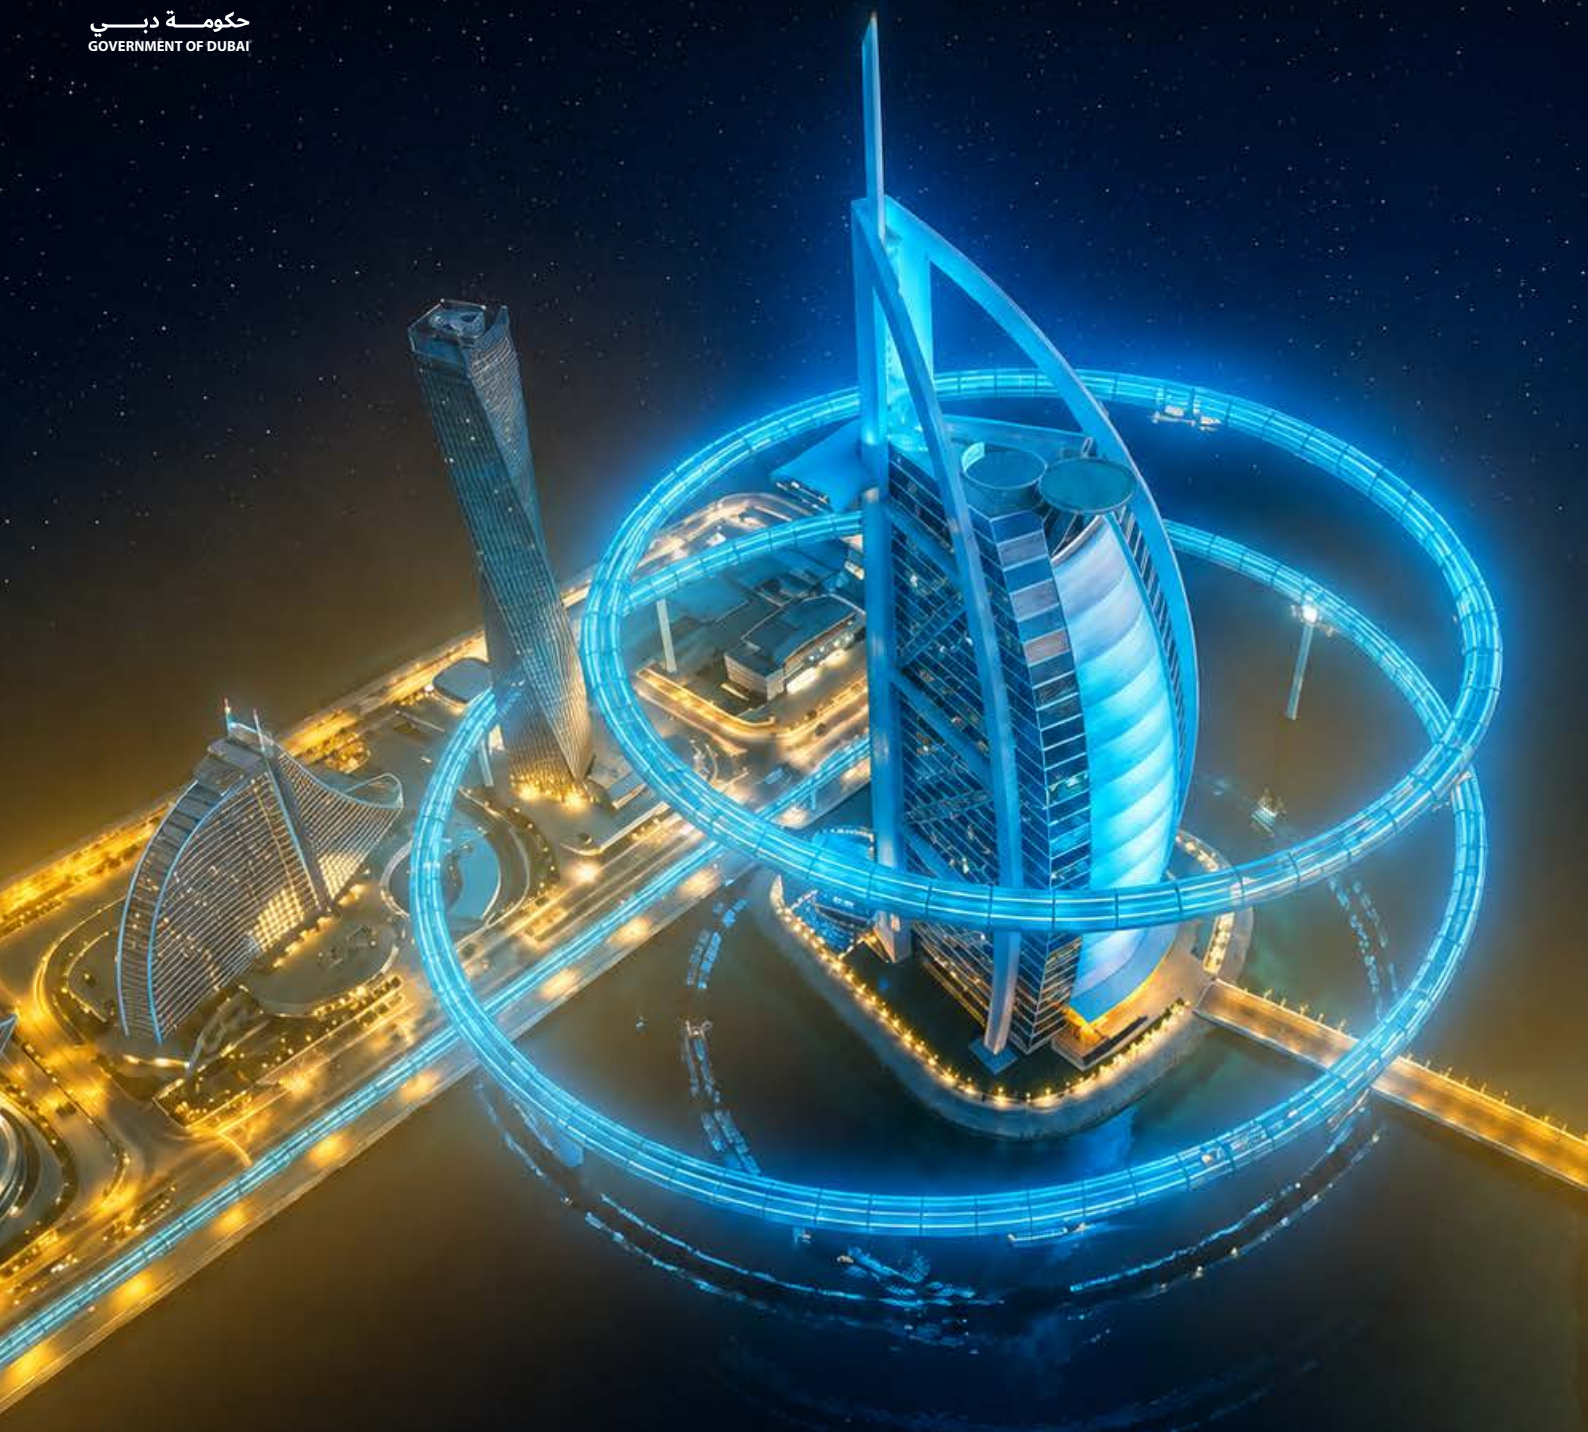
حكومة دبي
GOVERNMENT OF DUBAI

مركز دبي لاستخدامات الذكاء الاصطناعي
Dubai Centre for Artificial Intelligence

مؤسسة دبي للمستقبل
DUBAI FUTURE FOUNDATION

دبي الرقمية
DIGITAL DUBAI



Dubai's State of AI Report

Scaling with *Trust*

2026



The world is changing.
Technology is advancing and accelerating.
But one principle remains unchanged.
People come first. Our goal is to create the best environment to unlock their potential, improve their lives, and build a better government — a government that is faster, more responsive, and more impactful.

His Highness Sheikh Mohammed bin Rashid Al Maktoum
Vice President and Prime Minister of the UAE and Ruler of Dubai.

Statement on the UAE Government framework for agentic AI, April 2026 [1].

REFERENCE

[1] UAE Cabinet, "Mohammed bin Rashid reveals new UAE Government framework to deploy agentic AI across 50% of government sectors' operations within two years," uaecabinet.ae, Apr. 2026. [Online]. Available: https://uaecabinet.ae/en/news/under-directives-of-uae-president-and-in-world-first-mohammed-bin-rashid-reveals-new-uae-government-framework-to-deploy-agentic-ai-across-50-of-government-sectors-operations-within-two-years



Today, we are entering a new phase in this journey, with the ambition of building a world-leading government model in which AI becomes an integral part of how government operates and responds to people's needs."

His Highness Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum
Crown Prince of Dubai, Deputy Prime Minister and Minister of Defence of the UAE [1].

REFERENCE

[1] The National, "Sheikh Hamdan: 'Our goal is to make Dubai the world's leading AI hub'," thenationalnews.com, Jun. 11, 2026. [Online]. Available: https://www.thenationalnews.com/news/uae/2026/06/11/sheikh-hamdan-our-goal-is-to-make-dubai-the-worlds-leading-ai-hub/



FOREWORD

H.E. Omar Sultan Al Olama

Minister of State for Artificial Intelligence, Digital Economy, and Remote Work Applications

"The pace described here is possible only because people trust the institutions setting it. That trust took years to build.

In the past eighteen months, artificial intelligence has become central to how the UAE governs. Since January 2026, a national AI system has served as an advisory member of the Cabinet and the boards of federal entities and government departments. In April, under the directives of the UAE's leadership, the UAE Government set a national target: to deploy agentic AI across half of government sectors, services and operations within two years. In May, the Cabinet approved the framework to deliver this target, along with the largest training program in the Government's history, reaching 80,000 federal employees.

During the same period, the UAE has also expanded its efforts to prepare people for an AI-driven future. Artificial intelligence has been taught as a subject in government schools, from kindergarten to Grade 12, since last year. This September, the Cabinet extended the curriculum to every public and private school, with 22,000 teachers to be trained. National-scale computing capacity is also being developed within the UAE. Individually, these are significant decisions. Together, they signal a clear direction: to build, govern and use AI at national scale, while preparing people to make the most of it. The principle that opens this report sets the standard for all of these efforts: people come first.

National targets are ultimately delivered service by service, and this is where Dubai plays a critical role. Dubai serves millions of residents, businesses and visitors every day, and it is advancing these ambitions across government. Chief AI Officers now lead this work across their government entities. A common policy ensures that these entities operate according to the same principles, while a sovereign platform provides the infrastructure for running AI models. On 1 April 2026, Dubai set a one-year target to bring every government service onto a single AI-powered platform.

This report documents that progress, drawing on the experience of twenty government entities. It counts only systems that are live. It identifies what stalled as well as what worked, and gives outside critics a chapter of their own. For the national program, this provides practical evidence of what it takes to move AI from pilot to daily operation, the value it delivers, and what is required to govern it effectively.

The theme is well chosen. The pace described here is possible only because people trust the institutions leading this transformation. That trust took years to build. It came from systems that work and rules that are clear, and once lost, it is difficult to regain. The challenge now is to scale AI while maintaining that trust. Dubai has put its experience and results forward for others to examine. I commend this report to colleagues across the federal government and our sister emirates, to the partners who build with us, and to every government working towards the same goal.



FOREWORD

H.E. Hamad Obaid Al Mansoori

Director General, Digital Dubai Authority

"An intelligent city has one goal: a simpler life for the people who live in it.

Dubai has rebuilt the way its government reaches people three times in twenty-five years. In 2001 it put services online. In 2013 it moved them to the mobile phone. By the end of 2021 it had eliminated paper transactions across government. Each change used a different technology for the same purpose: to make life easier for the people and businesses of this city. That purpose comes from the vision of HH Sheikh Mohammed bin Rashid Al Maktoum, Vice President and Prime Minister of the UAE and Ruler of Dubai, and it has not changed. Artificial intelligence is the fourth change, and the largest.

For twenty-five years we improved services. The next step is to stop asking people to think in services at all. Someone who wants to open a restaurant in Dubai needs a trade licence, a food permit, a health and safety permit, a bank account and visas for staff. That is five separate requests for one need. Our direction is to move from providing services to meeting people's needs. Our ambition is to enable people to express their needs once, in their own words, and have the relevant services and processes coordinated across government and participating private-sector partners, with appropriate safeguards and human oversight.

Dubai is well placed to do this. It benefits from a unified strategic vision, decisive leadership, and a government ecosystem capable of aligning institutions around shared priorities. Government entities are building on shared digital infrastructure, common data standards, and a unified AI policy, creating stronger foundations for interoperability, collaboration and the reuse of solutions.. The results are being noticed. This year Dubai ranked second among 61 cities in Boston Consulting Group's Intelligent Cities Index, and first for the adoption of AI technologies and smart city solutions.

There is still a great deal to build, and no single organization can build it. Digital Dubai will work with all government entities and with our partners in industry and academia as one team, as we have at every earlier stage. We will also share what we learn. Cities everywhere are asking how to bring AI into daily life without losing people's trust, and Dubai intends to be the place they look to for a working answer. This report is where we account for our progress each year. Twenty government entities set out what they have deployed and what it changed. Outside contributors test that record, and the final chapter states the principles we will apply as we scale. Read it as a measure of how far Dubai has come, and return to it next year to see how far the city has moved from services to needs. An intelligent city has one goal: a simpler life for the people who live in it.



FOREWORD

H.E. Khalfan Belhoul

Chief Executive Officer, Dubai Future Foundation

"Foresight is not prediction. It is about making decisions today that remain sound across different possible futures.

Foresight is not prediction. It is about making decisions today that remain sound across different possible futures. That distinction matters more in 2026 than it did a year ago, because the choices governments make today on artificial intelligence will have consequences for years to come. Decisions on infrastructure, procurement and talent will continue to shape public services in 2030 and well beyond.

One major shift is the rise of agentic AI. AI systems are moving from answering questions to taking action, from producing a draft to executing a task, checking the result and deciding what to do next. This changes the challenge for government. The question is no longer simply what AI can do, but what we should delegate to it, how it should be supervised and what happens when something goes wrong. These are questions of institutional design, and institutions take time to adapt. We need to address them before these technologies become widely adopted, not after.

Two other shifts deserve attention. The first is talent. The skills we need are changing from simply building AI systems to knowing how to direct them: setting clear objectives and boundaries, and assessing their outcomes. A workforce prepared for the last technology cycle will not necessarily be prepared for this one. Retraining at the scale of government takes time, and it needs to begin early.
The second is the growing concentration of AI capabilities. Computing power, data and advanced AI capabilities are increasingly concentrated among a smaller number of players globally. This makes building sovereign capabilities increasingly important. Cities that invest in building their own capabilities, rather than relying only on what they can buy, will be better positioned for the decade ahead. These shifts are not theoretical; they are already visible in the submissions gathered for this report.

Dubai Future Foundation's contribution to this report is to bring a long-term perspective to the operational work taking place today. Both perspectives are essential. Decisions about scaling AI without considering what lies ahead can lead to costly corrections later, while foresight without a way to put it into practice has limited value. What gives this report its weight is that Dubai is doing both: implementing AI today while preparing for what comes next, and documenting what it learns along the way.

The principles set out in the final chapter should be read in that context. Their value is not in providing definitive answers to today's questions, but in helping government respond to change as AI continues to evolve. That is what preparedness looks like in practice: not preparing for one version of the future, but building the ability to respond to many.



FOREWORD

H.E. Younus Al Nasser

Chief Executive, Dubai Data and Statistics Establishment

"A city can only meet the needs it understands, and that understanding begins with data.

At the end of 2025, 4.58 million people lived in Dubai. In the busiest hours of the day the city holds 6.39 million on average, as people arrive to work, study, shop and visit. A city that planned only for its residents would overlook 1.8 million people every day. Behind every number in this city there is a person, and Dubai has set out to meet that person's needs. A city can only meet the needs it understands, and that understanding begins with data.

Artificial intelligence makes this more important. The same AI models are now within reach of every city. What differs is the data each city brings to them, and a model's answers are only as reliable as the data behind them. Strong AI requires strong data foundations. Dubai has been building those foundations for more than a decade, and it now manages its data as a strategic asset. Its data law was issued in 2015, and this year an updated Dubai Data Manual gave every government entity one framework for governing and sharing its data and preparing it for AI.

The city's new direction asks more of its data. When a person states a need once, relevant information should be securely accessible to authorised entities, subject to applicable data governance, privacy and sharing requirements. Our work at the Dubai Data and Statistics Establishment is to make Dubai's data trusted, connected and ready. Trusted means it is accurate and protected, and used lawfully and securely, with appropriate authorisation, oversight and safeguards.. Connected means it is kept to common standards, because few needs sit within one organization. Ready means it is current and complete enough for an AI agent to act on.

Good data also changes how government decides. Dubai has set a target for AI-enabled data and analytics to inform at least 80% of government policies and strategies. This report asks the same of AI itself: that it be judged on evidence. Twenty government entities across seven sectors reported against one template, so results can be compared. The 2027 edition will be read against this year's results, and will go a step further, from counting where AI is in use to measuring what it changes in people's lives.
Data on its own changes nothing. Its value is in what the city does with it: a government that sees a need early, responds to it as one, and can show afterwards that it made a difference. That is what these foundations are for. We will keep building them, and keep measuring, until the evidence shows life in Dubai getting simpler for the people behind the numbers.



FOREWORD

H.E. Matar Al Hemeiri

Chief Executive, Digital Dubai Government Establishment

Our ambition is for Dubai to have the most intelligent government in the world, and designed around people needs. We want a government so intelligent that people rarely need to think about it, and so human that it never feels like a machine. Artificial intelligence brings that ambition within reach. For the first time, technology can understand what a person needs and act on it, which means government can be designed around life rather than around process. Building that government is the work of the Digital Dubai Government Establishment in the years ahead.

> *"Our ambition is for Dubai to have the most intelligent government in the world, and designed around people needs.*

For people, that means a government that anticipates rather than waits. Services will arrive at the moment a life changes. AI agents will act on people's behalf, always with their consent, carrying a request through every step until it is done, and routine obligations will take care of themselves. People come first, and the technology works quietly behind them. What they gain is the most valuable thing a city can give back: less noise in their days, and more time for the lives they want to live.

For government, it means working as one. People's lives do not follow organizational charts, and their experience of government should not either. Entities will collaborate by design, with fewer unnecessary handovers, simpler processes and a more consistent experience across government.. AI agents will take on routine work, predictive operations will anticipate demand, and leaders will decide with live insight in hand. Government will become faster and more efficient, and its people will spend less time on paperwork and more on judgment, empathy and invention.

Trust is the condition for all of it. AI agents must operate within clear authorization frameworks, with appropriate transparency, human oversight and accountability for their actions and outcomes.. Ambitions of this kind are never achieved by one organization, so we will pursue them with every government entity and with the companies, universities and innovators who share this city with us. Many of the journeys that matter most in a person's life cross between government and business, and they should feel like one.

This report shows that the journey is well under way. Twenty government entities describe how artificial intelligence is already improving their services and operations, and what they have learned along the way. Partners from industry and academia test that record from the outside, and the report sets out what it will take to scale with trust. It is an honest account of progress and a clear view of the work ahead. We will measure our success by the difference we make in people's lives: simpler experiences, greater trust, more effective government and more time for what matters to them.

1 Scaling with Trust

The 2026 thesis: the question is no longer whether AI works in government. It is whether government can scale it without spending the trust it runs on.



The question of 2026 is not whether AI works in government. It is whether government can scale it without spending the trust it runs on.

Something changed between the last edition of this report and this one, and it is worth stating plainly. Artificial intelligence stopped being a tool that government workers use and became, in a growing number of systems, an actor that does government work. Software now executes multi-step tasks, checks its own output, and decides what to do next. The 2025 report asked where AI stood in Dubai's government. The evidence from the twenty participating entities answers it: AI has moved materially into production across many areas of Dubai Government.
The question of 2026 is harder. Can a government scale AI without spending the trust it runs on?

This report argues that scaling responsibly is now the central institutional challenge of the moment, ahead of capability, ahead of investment, ahead of talent, because it is the constraint that decides whether the other three compound or unravel.
A pilot that fails costs a budget line. A scaled system that fails costs public confidence, and confidence is the one input a government cannot procure.
Scale changes the stakes before it changes anything else.

The timing is not incidental. Two forces converged this year to make the question urgent. The first is technical: agentic systems collapse the distance between a model's judgment and a public outcome, removing the human checkpoints that made earlier deployments forgiving. The second is institutional: the world's regulatory poles are diverging rather than converging, which means no government can wait for international consensus to settle the rules. Every administration is now writing its own answer in real time, whether deliberately or by default. Dubai has chosen to write its answer deliberately, and to publish it.

It also changes the nature of the problem. At pilot scale, difficulties look technical and can be fixed by better engineering. At production scale, they resolve into governance tensions: pairs of public goods in which a gain on one side can come at the expense of the other unless the institution is designed to hold both. This report identifies five, and uses them as its analytical lens from first page to last. Speed against safety: every week saved in deployment is a week not spent hardening the system. Automation against human oversight: every task delegated to software is a judgment removed from a person. Innovation against regulation: every rule that prevents harm also narrows the space in which better answers are found. Performance against explainability: the most capable models are often the hardest to interrogate. Centralisation against autonomy: what the centre standardises, the entity can no longer adapt. None of these can be escaped. All of them can be governed. The difference between the two is doctrine.

> *"Tensions cannot be escaped, but they can be governed. The difference between the two is doctrine.*

Dubai's response to these tensions is the subject of this report, and it is best understood as an architecture rather than a program. Its coordinating institutions steer a government of autonomous entities without flattening them. Its governance instruments work by incentive as much as by mandate, making the trusted path the attractive path rather than the only one. Its measurement systems, built on structured submissions from twenty government entities across seven sectors, turn compliance from an assertion into a number. The architecture is not offered here as a finished object. It is offered as a working answer, one that the middle chapters of this report test honestly against the evidence, including the places where the evidence pushes back.

That evidence base is what separates this report from a position paper. Twenty entities answered the same structured template: investment, deployed applications, workforce, compliance. Seven sector analyses read the patterns across them. Nine external voices, six from industry and three from the academy, examine the argument from outside and were not asked to agree with it. The rule throughout is that only deployed systems count. Pilots, prototypes, and intentions, however promising, are not treated as achievements, because residents cannot feel an intention. What this report claims, it can source. What it cannot source, it does not claim.

The year's accumulated lessons resolve, in the final chapter, into Dubai's Emerging Doctrine for an AI-Enabled City: one governing objective, public value before technological novelty, and five operating principles, one for each tension. Move at mandate speed, with assurance proportionate to consequence. Automate execution, retain explicit accountability. Enable experimentation through clear instruments; regulate deployment according to risk. Pursue the strongest performance that remains accountable for the intended use. Centralize shared foundations; federate service judgment and accountability. They are written as obligations, not descriptions: each names what the city's government should expect of itself before a system is deployed or scaled, the evidence it should require, and the cost it accepts. A strategy tells an institution what to pursue. A doctrine tells it how to decide when pursuits collide. In a field that changes monthly, the second is worth more.

None of this argues that the tension between scale and trust has been resolved in Dubai. It argues something more useful: that the tension has been named, instrumented, and placed under management, and that the record of doing so is worth publishing, frictions included, because trust compounds only when reasoning is shown. The chapters that follow deliver on the claims made here. The world changed; Chapter 1 measures it. Dubai built an architecture; Chapter 2 explains it. The architecture meets reality; Chapter 3 provides the evidence. The argument meets its critics; Chapter 4 gives them the floor. And what survives becomes doctrine; Chapter 5 writes it down. That is the whole of the report's method: claim, evidence, scrutiny, principle.

Scaling with trust is not the finish line of Dubai's AI program. It is the operating condition of everything that comes after. This report is the account of a government that has accepted that condition, and organised itself accordingly.





2 How to Read This Report

A Note from the Editorial Team

TABLE OF CONTENTS

Forewords iii
Executive Summary viii
How to Read This Report x

CHAPTER 1
The Inflection 1

CHAPTER 2
The Architecture 17

CHAPTER 3
AI in Action 57

CHAPTER 4
The Stress Test 101

CHAPTER 5
The Emerging Doctrine 123

Reference & Contributors 141



Reader,

This report has one analytical lens, and it never changes. Five trade-offs run through every chapter that follows: speed against safety, automation against human oversight, innovation against regulation, performance against explainability, centralisation against autonomy. Wherever you open the report, you are looking at one or more of them. They are defined properly at the close of Chapter 1; everything afterwards puts them to work.

The structure follows an argument, not a catalogue. Chapter 1 establishes what has changed, globally, regionally, and at home. Chapter 2 describes the architecture Dubai built in response. Chapter 3 tests that architecture against evidence from twenty government entities. Chapter 4 invites scrutiny, including from voices outside government. Chapter 5 distils what the year taught into doctrine. Inflection, architecture, evidence, scrutiny, doctrine: each chapter earns the next.

Three editorial choices shape what you will find. We publish friction, not just progress; where systems stalled or data disappointed, the report says so. We name limitations, including the limits of our own evidence base. And we include external voices, academic and industry, whose conclusions we did not tune to our own. A report on trust that managed its critics would refute itself.

One caution the reader deserves: this report is a snapshot taken from a moving train. The field changes monthly; the doctrine in the final chapter is designed to outlast the details, but details will age. Read the principles as durable and the figures as dated the day after publication.

Our thanks go to the teams in all twenty participating entities who compiled their submissions with patience and candour; to the directors general and executives who gave their time for the sector reflections; to our academic and industry contributors; to the statistics teams who held the evidence to standard; and to the designers who made a dense report readable. Their names do not appear on these pages. Their work is on every one of them.

The Editorial Team, Dubai State of AI Report 2026



The Inflection

What changed, and why it changes everything that follows

Every edition of this report begins by taking bearings. This year, the bearings themselves have moved. Since the 2025 baseline, artificial intelligence has crossed from a technology that assists work to one that performs it, and the crossing happened at every altitude at once. Globally, compute, capital, and capability concentrated at a pace that reset assumptions about who can build what. Regionally, ambition turned into infrastructure. At home, execution replaced announcement as the measure of progress.

This chapter records those shifts in that order: world, region, country. It then examines the technologies driving them and the trends most likely to shape the next planning horizon. It closes with the analytical instrument the rest of the report depends on: five trade-offs that every scaling government must govern, whether it names them or not. Read them carefully. Every chapter that follows, from the architecture Dubai has built to the doctrine it now applies, is an answer to the questions this chapter raises.

1 The Global Inflection

Three forces moved together in the past year: compute, agency, and capital. Their convergence is what makes 2026 different in kind, not just in degree

$581.7B
global corporate AI investment, 2025
+130% year-over-year · +127% private investment [1]

$150B+
one hyperscaler's reported total annual capex, 2025 [2]



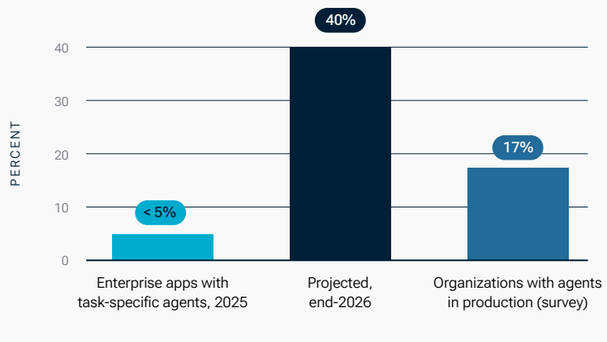


Figure 1.1: The agentic wave: projection versus production.

Inflections are easy to declare and hard to prove. The claim that the past year marks one rests on three forces that moved separately for a decade and, in 2025, began to move together: the industrialisation of compute, the shift of AI systems from answering to acting, and a reallocation of capital toward AI measured in the hundreds of billions of dollars. Each force alone would justify a government's attention. Their convergence demands it.

FIRST FORCE Compute becomes industrial policy

The first force is physical. Training and running frontier AI systems has become an exercise in heavy infrastructure: gigawatt-scale data centres, multi-year silicon supply agreements, and energy contracts that reshape national grids. The world's largest technology firms are now spending on infrastructure at a rate previously associated with national utilities; Google alone reported more than 150 billion US dollars total annual capital expenditure in 2025, amid sharply rising spending on AI infrastructure [2]. The consequence for governments is direct. Compute has joined energy and water on the list of strategic inputs, and access to it is decreasingly a market question and increasingly a diplomatic and industrial one. Independent trackers of training-compute growth, such as Epoch AI, point in the same direction: the frontier is an infrastructure race before it is a research one. Sovereign capacity, meaning compute built within one's own borders and jurisdiction, has accordingly moved from a theoretical debate to budget lines on several continents. The scarce resource in AI is no longer the idea. It is the electricity, the silicon, and the institutional capacity to absorb them.

> *"The scarce resource in AI is no longer the idea. It is the electricity, the silicon, and the institutional capacity to absorb them.*

SECOND FORCE The agentic wave

The second force is behavioural. Through 2025, AI systems crossed a threshold that separates a tool from an actor: they began to plan, use software tools, check results, and correct course across multi-step tasks without a person directing each step. Industry analysts expect task-specific agents to be embedded in some 40 per cent of enterprise applications by the end of 2026, up from under 5 per cent a year earlier [3]. Adoption surveys tell a more sober story: in one major survey of technology executives, only 17 per cent of organizations had agents in production, even as a large majority expected to deploy within two years [4]. Both numbers matter. Together they describe a technology arriving faster than the management practices needed to govern it, and a wave of institutional decisions that most organizations have not yet made. When software stops answering and starts acting, the question shifts from what a model knows to what an organization permits, supervises, and can reverse.

THIRD FORCE Capital reallocates

The third force is financial. Global corporate investment in AI, Stanford's aggregate of private investment, mergers and acquisitions, minority stakes and public offerings, reached 581.7 billion US dollars in 2025, an increase of 130 per cent over the prior year, with private investment alone growing by more than 127 per cent [1]. Behind the totals sits a quieter pattern: capital is concentrating by firm, by geography, and by layer of the technology stack. Investment at this scale pulls talent, suppliers, and research agendas along with it, which is why the reallocation is better read as a restructuring of the technology economy than as an enthusiasm. It is also unevenly distributed, and the unevenness compounds. Regions that secure infrastructure and talent early become the venues where the next round of capability is built; regions that wait become customers of it. For governments, the reallocation carries a second-order effect that outweighs the first: it sets the market price of everything a public AI program needs. Salaries for scarce profiles, terms for compute capacity, delivery schedules for hardware: all are now set by an investment cycle that governments do not control but must budget within. Public plans costed against last year's prices are already wrong.

What convergence changes

Any one of these forces could be managed within existing institutional habits. It is the interaction that resets the terms. Industrial compute makes capability a function of infrastructure; the agentic wave turns that capability into autonomous action inside real processes; concentrated capital decides where both accumulate. For a government, this collapses the comfortable distinction between technology policy and core statecraft. Decisions about data centres are now decisions about service delivery. Decisions about procurement are now decisions about supervision and recourse. And the pace is set externally: no administration controls the tempo of the frontier, and a widening gap has opened between what AI systems can do and how prepared institutions are to manage them [1].

> *"When software stops answering and starts acting, the question shifts from what a model knows to what an organization permits.*

This is the inflection against which the rest of this report is written. It is global, it is measurable, and it is indifferent to any single government's readiness. The two articles that follow trace how the same three forces are expressing themselves in this region and in the UAE, where the response has been not to slow the wave, but to build the institutions that can ride it. The five trade-offs that close this chapter are, in the end, the management form of the forces described here. The world changed first. The question is who changes next, and how deliberately.

REFERENCE

[1] Stanford Institute for Human-Centered Artificial Intelligence, The 2026 AI Index Report. Stanford, CA, USA: Stanford Univ., 2026. [Online]. Available: https://hai.stanford.edu/ai-index/2026-ai-index-report

[2] Stanford Institute for Human-Centered Artificial Intelligence, "Economy," in The 2026 AI Index Report, 2026. [Online]. Available: https://hai.stanford.edu/ai-index/2026-ai-index-report/economy

[3] Gartner, Inc., "Gartner predicts 40% of enterprise apps will feature task-specific AI agents by 2026, up from less than 5% in 2025," Press Release, Aug. 26, 2025. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025

[4] Gartner, Inc., 2026 CIO and Technology Executive Survey, 2025. [Online]. Available: https://www.gartner.com/en/articles/hype-cycle-for-agentic-ai



2 The Regional Inflection

The Gulf has stopped importing the AI era and started building it, a shift from adoption to capability that may be creating a third pole in the global system

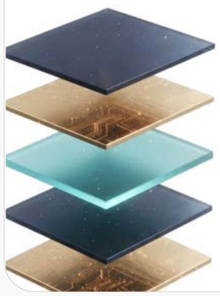
> *The region's defining move of the past year was not a purchase. It was a build.*

For most of the past decade, the Gulf's relationship with advanced technology followed a familiar sequence: observe, procure, adopt. The past eighteen months broke the sequence. Across the GCC, the defining move of the period was not a purchase but a build: of compute, of companies, of talent pipelines, and of the institutional machinery that turns all three into durable capability. The region has stopped importing the AI era and started constructing its share of it.

From adoption to capability

The evidence is concrete and recent. Sovereign AI ventures wholly owned by national investment institutions have been established to keep data, processing, and intellectual property within national jurisdiction [1]. Multi-gigawatt AI campuses are under construction, developed with the participation of the leading global technology firms; the first phases are scheduled to come online in 2026 [2]. Dedicated AI investment vehicles measured in the tens of billions of dollars have been capitalised from the region [3]. Hyperscale cloud regions designed for AI-intensive workloads are being built onshore [1]. Each project differs in scale and structure. The pattern does not: in every case, the objective has shifted from consuming AI services produced elsewhere to owning the productive capacity itself. Nor is the build only physical. National AI companies are assembling research and engineering teams onshore. Universities across the region have established dedicated AI faculties and graduate programs. Regulatory and governance functions that did not exist three years ago now employ specialised staff. Taken together, the direction of travel is toward a full stack: energy at the bottom, applications at the top, and, for the first time, every layer between them represented within the region rather than rented from outside it.

Every layer of the stack is now built within the region, energy to applications.



Why sovereignty, why now



The turn to sovereign capability is sometimes read abroad as prestige spending. The more accurate reading is insurance. The previous article described a world in which compute is concentrating, capital is consolidating, and access to frontier capability increasingly depends on infrastructure and jurisdiction. In such a world, a region that runs government services, financial systems, and critical infrastructure on AI cannot prudently leave every layer of that stack outside its own borders. Sovereign capacity is how the region converts exposure into position. It also compounds: infrastructure attracts talent, talent attracts research, research attracts the next round of capital. The Gulf's advantages here are real: energy, land, capital, and the ability to move from decision to construction quickly. The past year demonstrated all four at once, on schedules that other regions plan in decades and the Gulf now contracts in quarters.

A third pole

> *Between a regulatory pole and a capability pole, the Gulf is assembling something distinct: capability under governance.*

The global AI system is often described as bipolar: one pole organised around regulatory reach, whose instruments are conformity assessment and market access; another organised around frontier capability, whose instruments are compute, capital, and research depth. The past year supports an emerging proposition: that a third pole is forming in the Gulf, one that is neither a regulatory bloc nor a frontier-research cluster, but a deployment pole: a region that builds serious infrastructure, adopts at government scale, and governs through incentive and coordination rather than through either laissez-faire or heavy conformity regimes. Capability under governance is the shortest description. The claim is not that this pole rivals the other two in research output. It is that it offers something the other two do not: a working demonstration of what state-scale adoption looks like when infrastructure, governance, and deployment are planned together. For the many states that are neither regulatory superpowers nor frontier-research economies, which is to say most states, that demonstration may prove the most consequential export the region produces this decade.

The gap that remains

Honesty requires the counterweight. Infrastructure is the fastest layer of capability to build; the slower layers are talent depth, research originality, and the operational experience that only accumulates through years of running systems in production. Announcements are also not deliveries: construction schedules, energy commitments, and chip supply all carry execution risk, and independent analysts have noted the distance between declared ambition and commissioned capacity across the region [4]. Nor does infrastructure guarantee demand: compute built ahead of the applications that would use it is a bet on adoption, and the bet pays off only where governments and firms actually deploy at scale. The regional inflection is real, but it is a beginning. What converts a building program into a durable pole is the unglamorous work that follows: operating standards, measurement disciplines, workforce pipelines, and governance that earns trust at scale.

That work is where the UAE has moved decisively, from national frameworks to operating practice. The next article turns to it. The regional story explains why capability is being built here. The national story explains what it takes to make capability govern.

REFERENCE

[1] Middle East Institute, "AI, the Gulf, and the US: A primer," Washington, DC, USA, 2025. [Online]. Available: https://mei.edu/report/ai-the-gulf-and-the-us-a-primer/

[2] G42, "Global tech alliance launches Stargate UAE," Abu Dhabi, UAE, 2025. [Online]. Available: https://www.g42.ai/resources/news/global-tech-alliance-launches-stargate-uae

[3] Middle East Institute, "From crude to compute: Building the GCC AI stack," Washington, DC, USA, 2025. [Online]. Available: https://mei.edu/report/from-crude-to-compute-building-the-gcc-ai-stack/

[4] Gulf International Forum, "Gulf AI infrastructure: Examining the business and economic case," 2025. [Online]. Available: https://gulfif.org/gulf-ai-infrastructure-examining-the-business-and-economic-case/



3 The UAE Inflection

In one year the national posture moved from strategy to execution, and this report holds itself to the same standard, by measuring 2026 against what 2025 promised

> *The honest measure of a strategy is not its launch. It is what remains standing a year later and what had to change.*

Last year's edition of this report ended on commitments. This one begins by checking them. Between the 2025 baseline and today, the UAE's AI posture moved from strategy to execution: visibly, at the federal level, and in ways a reader can verify from public record. The honest measure of a strategy is not its launch. It is what remains standing a year later, what accelerated, and what had to change. This article accounts for both columns: what evolved as planned, and what recalibrated on contact with reality. Both columns belong in the record, because a report that publishes only the first is a brochure.

50%

of government sectors' operations mandated for agentic AI deployment within two years

UAE Cabinet, April 2026 [1]

Autonomy is being extended task by task, with instrumentation,

rather than process by process on a schedule.



What evolved

The clearest single marker of the national inflection is federal and public. In April 2026, the UAE Cabinet approved a government framework to deploy agentic AI across half of government sectors' operations within two years, described by the leadership as a shift in which AI becomes an executive partner that analyses, decides, and executes in real time [1]. Whatever else that framework becomes, its existence settles a definitional question: the era in which government AI meant chat assistants and document summarisation is over at the highest level of national intent. Execution, not experimentation, is now the declared unit of progress.

The machinery of execution has also thickened below the federal layer. Chief AI Officers have been appointed across Dubai Government entities, giving each entity a named, accountable owner for AI adoption and a direct line into the coordination structures this report describes in Chapter 2 [2].

It is tempting to read such appointments as administrative housekeeping. The opposite is true. Strategies announce intent; appointments assign accountability, and the year in which every entity acquired an accountable owner is the year adoption stopped being optional.

The same movement from concept to operation is visible in the instruments this report examined a year ago. The AI Seal, introduced in the 2025 edition as an incentive-based trust mark, has moved from concept into enforcement practice, exactly the trajectory this report argued an incentive instrument should follow. Infrastructure aggregation, announced last year as intent, is now operational. And the coordination mandate at the centre of the governance architecture has accumulated a year of coordination practice. Chapter 2 examines each instrument in depth; the point here is the pattern. Across every instrument the 2025 report described in the future tense, the 2026 report can write in the present tense. The delta cards on this spread carry the specifics.

What recalibrated

A report that recorded only progress would not deserve the trust this one asks for. Three expectations were adjusted by experience, and each adjustment carries a lesson. First, agentic timelines. The ambition is federal and public; the practice of the past year taught that autonomous execution inside live government processes requires supervision design, exception handling, and rollback paths that pilots never surface. Industry experience points in the same direction: analysts project that a large share of early agentic projects will be cancelled before 2028 for want of clear value or adequate controls [3]. The recalibration is not retreat; it is sequencing. Autonomy is being extended task by task, with instrumentation, rather than process by process on a schedule.

Second, talent supply. The assumption that hiring could satisfy AI workforce demand met a global market in which every government and firm is bidding for the same profiles. The observed response has been a pivot from recruitment to conversion: upskilling serving employees into AI-adjacent roles, and redefining roles around directing AI systems rather than competing with them. Third, time-to-production. The distance between a working demonstration and a deployed service, once procurement, integration, assurance, and accountability sign-off are counted, proved consistently longer than the demonstration culture suggests. The lesson has been absorbed structurally: deployment-readiness is now treated as a property to be engineered from the start, not a phase appended to the end. That single change of posture may be worth more than any individual system shipped this year.

> *A government confident enough to publish its plans should be confident enough to publish its recalibrations.*

Accountability, then forward

This section exists because the 2025 report made claims about the future, and a government confident enough to publish its plans should be confident enough to publish its recalibrations. The pattern of the year is favourable but not tidy: national intent sharpened, instruments hardened into operation, and three assumptions were corrected in public. That is what execution looks like from inside. The remaining articles of this chapter widen the lens one final time, to the technologies forcing the pace and the trends that will shape the next horizon, before the report turns, in Chapter 2, to the architecture Dubai has built to hold all of it together.

THE AI SEAL

2025	2026
Incentive-based trust mark, newly introduced	**325** companies applied — representing 77 international offices; six tiers, E to S; certification required to partner with government on AI.

SOVEREIGN COMPUTE

2025	2026
Infrastructure aggregation announced as intent	**42%** average monthly growth — in the running total of tokens processed on llmapi.govai.ae since September 2025; 31 entities, 19 models.

ACCOUNTABLE AI LEADERSHIP

2025	2026
22 Chief AI Officers appointed	**27** entities' CAIOs — convened as a standing community, with the Dubai AI Acceleration Taskforce alongside it.

WORKFORCE CONVERSION

2025	2026
240 trainees in eight batches.	**1,000+** employees — across 78 entities, on a path to 50,000 by 2030.

REFERENCE

[1] UAE Cabinet, "Mohammed bin Rashid reveals new UAE Government framework to deploy agentic AI across 50% of government sectors' operations within two years," Apr. 2026. [Online]. Available: https://uaecabinet.ae/en/news/under-directives-of-uae-president-and-in-world-first-mohammed-bin-rashid-reveals-new-uae-government-framework-to-deploy-agentic-ai-across-50-of-government-sectors-operations-within-two-years

[2] Government of Dubai Protocol, "Hamdan bin Mohammed appoints 22 Chief AI Officers across government entities in Dubai," 2025. [Online]. Available: https://www.protocol.dubai.ae/en/media-listing/news-events/hamdan-bin-mohammed-appoints-22-chief-ai-officers-across-government-entities-in-dubai/

[3] Gartner, Inc., "Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027," Press Release, Jun. 25, 2025. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

4 Technology Updates

Three technical threads changed what is possible since 2025, and each one arrives in government not as a feature, but as an institutional decision

AGENTIC SYSTEMS

From answers to actions

The defining technical shift of the period is the maturing of agentic behaviour: systems that plan a task across multiple steps, call software tools, evaluate their own intermediate results, and decide what to do next. The mechanics matter less than the consequence: the human moved.
In assistant-era deployments, a person sat between every model output and every real-world effect. In agentic deployments, the person moves to the boundary: setting objectives, defining limits, auditing outcomes. Analysts expect task-specific agents to feature in a large share of enterprise software by the end of 2026 [1], which means governments will increasingly find agency arriving inside the tools they already buy, whether they asked for it or not.
The institutional implication is supervision design. An agent embedded in a licensing process is not a productivity feature; it is a delegation of administrative action, and it demands the same instruments any delegation demands: clear mandate, bounded discretion, logged decisions, and a route of appeal. This thread sharpens the speed-versus-safety and automation-versus-oversight trade-offs defined at the close of this chapter.

> *A technology becomes a government's problem, and its opportunity, at the moment it changes who decides.*

MULTI-MODAL

The interface dissolves

The second thread is quieter but broader. Systems that handle text, vision, voice, and documents in a single flow have moved from demonstration to dependable, and their effect on government is best understood through the humble form. A resident photographs a damaged pavement; the image is the report. A scanned contract, a voice note in either of the city's working languages, a decades-old paper record: each becomes directly processable without a human transcription step. The institutional implication is reach.
Services built on multi-modal systems meet residents in whatever medium the resident already holds, which collapses one of the oldest barriers between government and the governed: the form that had to be filled correctly. The same capability, pointed inward, opens archives that were never digitised beyond scanning: decades of paper records becoming, for the first time, queryable administrative memory. The reach runs through language too: systems that move between Arabic and English, in speech and in writing, remove a translation layer that has quietly shaped who finds government easy to use. The decision this thread forces is about quality assurance. When the input can be anything, validation can no longer live in the input form; it must move into the system itself, and be tested there.



Technology sections in government reports age badly, so this one follows a rule: every development described here is included not for its novelty but for the institutional decision it forces. Three threads met that bar since the 2025 edition: agentic systems, multi-modal capability, and the rise of smaller, specialised models. Each changes what is possible. More importantly for this report, each changes who decides, who supervises, and who answers when something goes wrong.

SMALLER, SPECIALISED MODELS

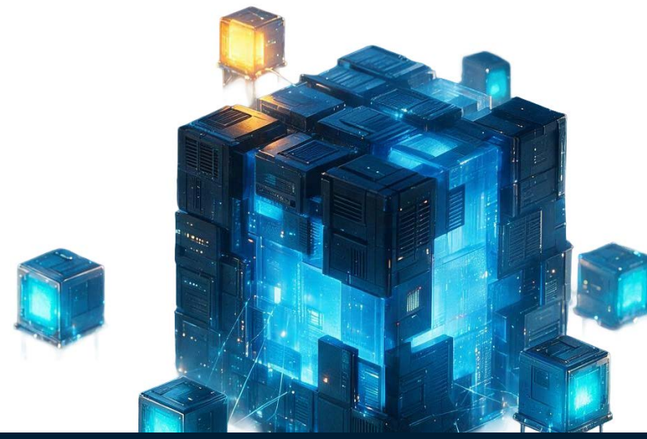
Fit-for-purpose

The third thread reverses an assumption that shaped early AI strategies: that capability would belong to whoever accessed the single largest model. The past year consolidated the opposite pattern, and it is one that favours disciplined buyers over deep-pocketed ones: families of smaller models, tuned to narrow tasks, running at a fraction of the cost, latency, and infrastructure demand of frontier systems, and in many production settings matching them on the task that matters.
For governments the implications stack. Economics: fit-for-purpose models turn AI from a metered external dependency into a deployable asset. Sovereignty: smaller models can run inside national infrastructure and jurisdiction. Explainability: a narrow model doing one job is easier to test, monitor, and defend than a general one doing many.
The decision it forces is portfolio discipline. The question is no longer which model is best, but which model is appropriate: for this task, at this risk, at this cost. That question is procurement's to ask, and this year it became answerable.

> *The question is no longer which model is best. It is which model is appropriate, for this task, at this risk.*

The common thread

Read together, the three threads describe a technology growing more capable and more governable at the same time: agency that can be bounded, interfaces that can be widened, models that can be right-sized. None of it is self-governing. Every gain described here is conditional on institutions doing deliberate work: supervision design, validation architecture, portfolio discipline. That conditionality is the bridge to what follows. It also carries a budget truth worth stating plainly: technology choices are now architecture choices, and architecture choices are governance choices; the three can no longer be made by three different committees on three different timelines. The trends article asks where these threads run next; the trade-offs essay turns them into the decision framework the rest of this report applies.



REFERENCE

[1] Gartner, Inc., "Gartner predicts 40% of enterprise apps will feature task-specific AI agents by 2026, up from less than 5% in 2025," Press Release, Aug. 26, 2025. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025

[2] Stanford Institute for Human-Centered Artificial Intelligence, The 2026 AI Index Report. Stanford, CA, USA: Stanford Univ., 2026. [Online]. Available: https://hai.stanford.edu/ai-index/2026-ai-index-report



5 Foresight & Trends

Not predictions but structured trends a government planner can act on, and the institutional posture each one demands

Prediction is a poor tool for governments: it concentrates risk on one future and expires on contact with events. This article offers something more modest and more useful: three trend clusters visible in present evidence, each stated with the institutional implication it carries. A trend is only useful to a government if it survives being wrong in the details. These are framed to survive.

> *"A trend is only useful to a government if it survives being wrong in the details.*

TREND ONE

Agentic operations become the default paradigm

The direction of enterprise software is now set: agency is being built into the tools organizations already use, and analysts expect it to be pervasive in application portfolios within the planning horizon [1]. The same analysts expect a large share of early agentic projects to fail or be cancelled [2]. Both projections can be true, and together they define the trend: planners should prepare for agentic operations to become the default way work is executed, and the differentiator will not be access to agents but the discipline with which they are governed. The institutional implication is to build the supervision layer before the agent estate grows: mandates, logging, exception routes, and kill-switches as standing infrastructure rather than per-project improvisations. Organizations that treat governance as the product of this transition, not the tax on it, will be the ones whose deployments survive the cancellation wave the analysts foresee, and inherit the ground vacated by those that do not.

TREND TWO

Regulatory divergence, not convergence

The comfortable assumption that global AI rules would converge on a common framework weakened further this year. The major jurisdictions are moving at different speeds, with different instruments, toward different objectives: conformity regimes in one pole, capability races in another, and incentive-based models emerging between them. Planners should assume a decade of divergence. The institutional implication is portability: systems, contracts, and assurance processes designed to satisfy the strictest plausible regime cheaply, rather than optimised for any single one. In practice that means model documentation kept to the most demanding standard in use, audit trails built in by default, and exit clauses in every dependency that crosses a jurisdiction. For a city whose economy runs on international flows, interoperability across regimes is not a compliance burden. It is a competitive asset, and it is buildable in advance.

TREND THREE

The AI-native workforce arrives

The third trend is demographic, and already under way. Within the planning horizon, a generation that has never worked without AI enters the workforce, while the serving workforce splits into those who direct AI systems and those who are displaced around them. The durable skill of the coming period is not operating AI; it is knowing what to delegate to it, what to withhold, and how to audit what comes back. The institutional implication runs deeper than training programs. Role definitions, career paths, performance measures, and management practice all assume that people do the tasks. That assumption is now false, and the institutions that redesign around direction and judgment, rather than task execution, will hold their talent while others bid for it. The redesign cannot wait for the new generation to arrive; by then, the institutions that started early will already employ it.

1 Agentic operations: the probable default

EVIDENCE

Agency is being built into the tools organizations already use, and analysts expect it to be pervasive in application portfolios within the planning horizon, while a large share of early agentic projects are expected to fail or be cancelled [1], [2].

IMPLICATION

Build the supervision layer before the agent estate grows, with mandates, logging, exception routes and kill-switches as standing infrastructure.

2 Regulatory divergence, not convergence

EVIDENCE

Major jurisdictions are moving at different speeds with different instruments, conformity regimes in one pole and capability races in another.

IMPLICATION

Design for portability, with documentation kept to the strictest plausible standard, audit trails by default and exit clauses in every dependency that crosses a jurisdiction.

3 The AI-native workforce arrives

EVIDENCE

Within the planning horizon a generation that has never worked without AI enters the workforce, while serving roles split between those who direct AI systems and those displaced around them.

IMPLICATION

Redesign roles, career paths and performance measures around direction and judgment rather than task execution.

> *"The durable skill of the coming period is not operating AI. It is knowing what to delegate to it.*

Why this matters for what follows

Three trends, one common property: none can be met with a fixed plan, because each will change shape en route. What they can be met with is a posture: principles stable enough to hold while the details move, and honest enough to revise when the evidence demands it. That is the case for the doctrine this report builds toward in its final chapter, and it is why the doctrine is written as operating principles rather than targets. The next article supplies the analytical frame those principles answer to. The trends say the ground will keep moving. The trade-offs say what must be balanced while it does.

REFERENCE

[1] Gartner, Inc., "Gartner predicts 40% of enterprise apps will feature task-specific AI agents by 2026, up from less than 5% in 2025," Press Release, Aug. 26, 2025. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025

[2] Gartner, Inc., "Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027," Press Release, Jun. 25, 2025. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

6

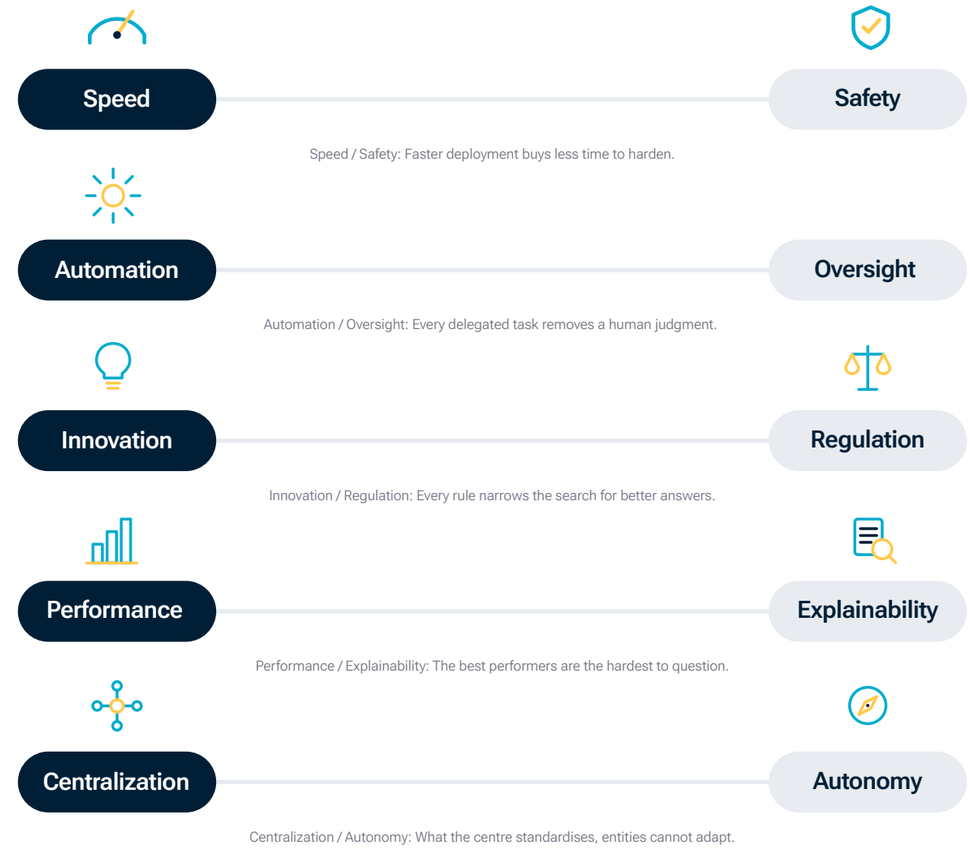
The Five Tensions of Scaling

Every government that scales AI must balance five tensions it cannot escape. This report names them, defines them, and uses them on every page that follows.

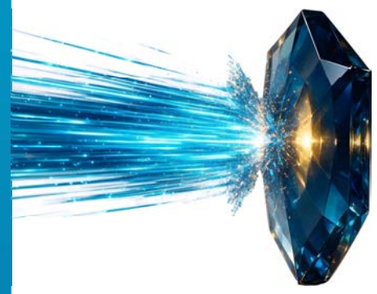


> *A trade-off differs from a problem in one decisive way: a problem can be solved, a trade-off can only be governed.*

Every discipline that matures acquires an analytical instrument: economics its supply and demand, medicine its differential diagnosis, engineering its factor of safety. The governance of AI at scale is a young discipline, and this report proposes its instrument: five tensions that every scaling government must balance, whether or not it ever names them. A tension differs from a problem in one decisive way: a problem can be solved and closed, while a tension can only be governed, continuously. Nor is it always zero-sum: poor institutional design forces a hard choice between the two goods, and good design can sometimes relax it, but no design removes it. The five below were selected against three tests. Each is real: both sides carry legitimate value, and while good design can ease the tension, none dissolves it. Each is consequential: mismanaging it produces failures measured in public trust, not merely in budgets. And each is decidable: a government can take deliberate positions on it, which means it belongs in doctrine rather than in commentary. Institutions that pretend a trade-off away do not escape it. They meet it later, unprepared, in public.



1 Speed versus Safety



DEFINITION

Every reduction in the time from capability to deployment is purchased with a reduction in the time available to harden, test, and rehearse failure.

EXAMPLE

Any live service: an agentic approval system shipped in weeks begins returning value in weeks, and carries into production the edge cases that a longer assurance cycle would have caught.

TENSION & COST

The tension is genuine because both sides compound. Speed compounds through earlier learning, earlier savings, and the operational experience that only production provides; safety compounds through failures that never happen. The cost of getting it wrong is asymmetric and delayed. Err toward safety and the loss is invisible: services that stayed slower for longer than they needed to, and a culture that learns to treat caution as achievement. Err toward speed and the loss is a public failure that arrives with names, dates, and consequences, and that taxes every subsequent deployment through restored approval layers. A scaled failure does not subtract one service; it subtracts permission.

2 Automation versus Human Oversight

DEFINITION

Every task delegated to a machine is a judgment removed from a person, including the judgments nobody realised the person was making.

EXAMPLE

Any casework queue. An automated eligibility system clears thousands of routine cases without fatigue, and the human who once processed them also noticed, without being asked to, the application that was technically complete but obviously wrong. Automate the processing and that unpriced inspection disappears with it.

TENSION & COST

The tension is genuine because human oversight is simultaneously the safeguard and the bottleneck: keep a person on every decision and AI becomes an expensive suggestion box; remove people entirely and the system's errors travel at machine speed with no one positioned to interrupt them. The cost of getting it wrong compounds quietly. Over-automation surfaces as unjust individual outcomes: each one small, each one accumulating into a reputation. Over-oversight surfaces as automation theatre: the cost of AI plus the cost of checking everything it does, with a supervising human whose attention has long since eroded into ritual approval. The design question is never whether humans oversee, but where: at the case, at the pattern, or at the boundary.

> *Institutions that pretend a trade-off away do not escape it. They just meet it later, unprepared, in public.*

3 Innovation versus Regulation

DEFINITION

Every rule that prevents a harm also narrows the search space in which better answers would have been found.

EXAMPLE

The sandbox problem in any jurisdiction: an entity with a novel use of AI needs room to try it; the regulator's duty is to imagine everything that could go wrong first. Wherever the line is set, it prices risk for every innovator who follows.



TENSION & COST

The tension is genuine because the two goods operate on different clocks. Innovation is fast, plural, and error-embracing; regulation is deliberate, uniform, and error-averse, and each is correct by its own logic. The cost of getting it wrong is measured in geography and in shadows. Over-regulate and innovation relocates to friendlier jurisdictions, taking talent and investment with it, while internal teams learn to stop proposing. Under-regulate and the harms land on the least protected, followed by emergency rules that overcorrect. There is also a subtler failure: regulation written for the last technology, faithfully enforced against the next one, achieving neither innovation nor protection. The instrument matters as much as the line: mandates forbid, incentives steer, and the difference between them is a recurring subject of this report.

4 Performance versus Explainability

DEFINITION

The systems that perform best on a task are frequently the hardest to interrogate about any single decision, and gains in one direction tend to cost the other.

EXAMPLE

The fraud-detection choice every large administration now faces: a complex model catches meaningfully more fraud than the transparent rules it replaces, but when a resident asks why their transaction was flagged, the honest answer is a probability, not a reason.

TENSION & COST

The tension is genuine because both properties are duties for a government rather than preferences. Performance is a duty because a worse model, knowingly deployed, wastes public money and misses harms it could have caught. Explainability is a duty because administrative decisions must be contestable: a decision that cannot be explained cannot be meaningfully appealed. The cost of getting it wrong: choose opacity and the first contested case in which the government cannot account for its own decision becomes a precedent that poisons trust in every system after it. Choose transparency at any price and the measurable cost is the fraud not caught and the need not met by the weaker model, a cost that is real but never makes headlines. The resolution is rarely a single point; it is a portfolio, with explainability priced by the severity of the decision.

> *The five tensions are this report's instrument panel. Every chapter that follows reads from it.*

5 Centralization versus Autonomy

DEFINITION

Everything the centre standardises is something the individual entity can no longer adapt; everything left to entities is something the government cannot guarantee uniformly.

EXAMPLE

Shared platform decision. Centralize a capability and every entity inherits security, purchasing power, and interoperability, along with a single point of failure and a queue for changes that fit no one's edge case perfectly. Leave the capability to entities and each one fits its own reality, and the government wakes to a dozen incompatible systems, a dozen security postures, and no common answer to a resident who reasonably expects one government, not a federation of exceptions.

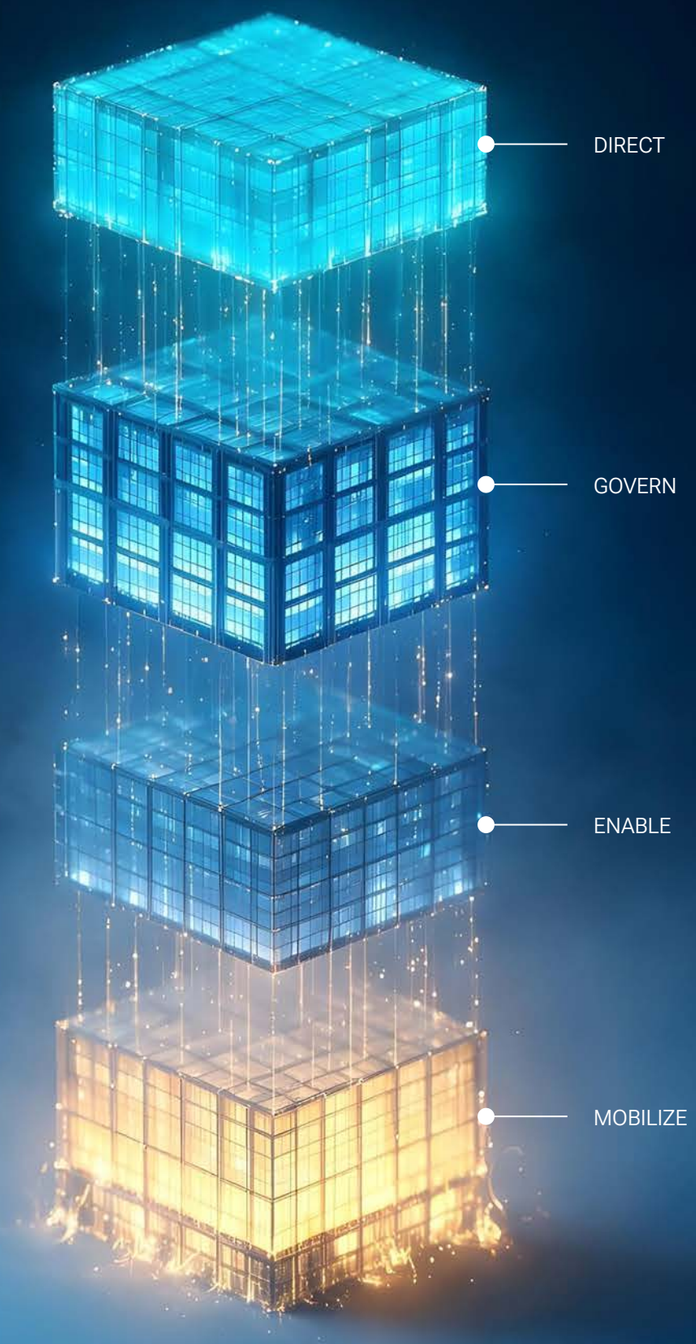


TENSION & COST

The tension is genuine because scale economies and local fit are both real, and because the failure of each pole is the argument for the other. The cost of getting it wrong is structural: over-centralize and innovation slows to the speed of the slowest committee while entities quietly build workarounds; over-devolve and every duplicated system is money spent twice and a boundary where accountability gets lost. Of the five tensions, this one is the most distinctly governmental: private firms can simply choose a structure, while a government of sovereign-mandated entities must engineer its balance deliberately and keep engineering it as the technology moves.

The instrument panel

Five tensions, one discipline. They interact: speed pressures oversight, centralisation reshapes regulation, explainability constrains automation. No static setting balances all five at once, which is why this report treats them as an instrument panel to be read continuously rather than a checklist to be completed. Nothing in this chapter claims a resolution, and the reader should distrust any document that does. What can be claimed is this: the tensions can be named, positions on them can be taken deliberately, and the taking can be written down and tested. Naming them in public is itself a governance act: it converts private engineering dilemmas into decisions that leadership owns, and it gives every entity in a federated government the same vocabulary for escalating them. The chapters that follow do exactly that. Chapter 2 describes the architecture Dubai built with these five tensions in view. Chapter 3 tests it against the evidence of twenty entities. Chapter 4 lets outside voices probe it. And Chapter 5 states the positions themselves: five principles, one per tension, each with its accepted cost. The backbone is now in place. Everything after this page is load.



The Architecture

Four layers, one operating model: how Dubai directs, governs, enables and mobilizes its way to AI at city scale

On 1 April 2026, HH Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai, Deputy Prime Minister and Minister of Defence of the UAE, and Chairman of The Executive Council of Dubai, gave the emirate a deadline: every government service on one AI-powered platform within a year. This chapter shows the machinery built to meet it, organized as the Dubai AI Operating Model, four layers, each answering one question.

DIRECT asks who sets the course: the Dubai Centre for Artificial Intelligence, DCAI (2.1) and the unified platform mandate (2.3). GOVERN asks what the rules are: the AI Policy (2.4), the AI Integration Matrix Framework (2.5) and the Federated Agentic AI Framework (2.11). ENABLE asks what the means are: CODE.AI (2.7), the AI+ workforce program (2.8), Infraboost compute (2.9), GRP agents (2.10) and AI Connect (2.12). MOBILIZE asks who builds with us: the Dubai AI Seal (2.2) and the private sector spotlights (2.13 to 2.15). The entities' own deployments fill Chapter 3. One page stands outside the model as its evidence: DDA by the Numbers (2.6).

The Dubai Future Foundation's DCAI mobilizes the city; the Digital Dubai Authority engineers and governs the architecture itself. Governance opens and closes the chapter: scaling with trust.

REFERENCE

[1] Dubai Media Office, "Hamdan bin Mohammed directs the integration of all government services into a single unified AI-powered digital platform," mediaoffice.ae, 1 Apr. 2026.

1 DUBAI CENTRE FOR ARTIFICIAL INTELLIGENCE:

Coordination Without Control

The institution that aligned an entire city's AI effort owns no entity's use cases, and that restraint is the design



Every government that takes artificial intelligence seriously eventually meets the same structural problem: the technology is horizontal, and government is vertical. Health data sits with one entity, transport telemetry with another, licensing records with a third, and the capabilities that could connect them belong to no one in particular. Cities usually respond in one of two ways. Some centralize, creating a single AI authority that executes on behalf of everyone and gradually crowds out the initiative of the entities it was built to serve. Others fragment, letting every department procure and govern alone, and discover years later that they have built the same capability many times over. Dubai chose a third path and gave it an institution. The Dubai Centre for Artificial Intelligence (DCAI), established in June 2023 as an initiative of the Dubai Future Foundation, coordinates the city's AI effort without controlling it.

> *"Twenty-seven Chief AI Officers, one taskforce, one blueprint, one annual week: DCAI's instruments align the city without taking a single use case away from the entities that own them.*

The convening power

DCAI's first instrument is a table with 27 chairs around it. The Dubai Universal Blueprint for Artificial Intelligence, unveiled in April 2024 under the leadership of HH Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai, Deputy Prime Minister and Minister of Defence of the UAE, and Chairman of The Executive Council of Dubai, required every government entity to appoint a Chief AI Officer, placing accountable AI leadership inside each organization rather than above it. DCAI convened the Chief AI Officers of 27 government entities into a single working community: a standing forum where the executives who own AI delivery inside their entities compare approaches, surface shared obstacles and align on citywide priorities. The distinction matters. The CAIOs do not report to DCAI; they report to their own leadership. What DCAI provides is the connective tissue that turns 27 parallel mandates into one conversation, so that a solution proven in one entity travels to the next as a matter of course rather than coincidence.

The Taskforce accelerates the work of others;
it does not take the work over.



Acceleration by taskforce

Convening alone does not create speed, so DCAI added a second instrument: the Dubai AI Acceleration Taskforce, which brings government and ecosystem leadership together to fast-track the adoption of AI across the city's priority areas. Where the CAIO community aligns direction, the Taskforce compresses timelines, identifying where an initiative is ready to scale and clearing the institutional path in front of it: the approvals, the interdependencies, the questions that cut across more than one entity's remit and would otherwise wait for an owner. Here again the centre stops short of execution. The Taskforce accelerates the work of others; it does not take the work over. Delivery stays with the entity that owns the service, the data and the accountability for both, which is why the acceleration holds: an entity moves fastest on work it still recognizes as its own.

DUBAI AI WEEK
2025 EDITION

250+
sessions

30+
launches

The city's AI stage

The third instrument is rhythm. Dubai AI Week, staged by DCAI, gathers the global AI community in the city for a concentrated week of policy, product and talent: the 2025 edition ran more than 250 sessions and saw more than 30 launches across its program. For the government audience, the week functions as an annual forcing mechanism, a date on the calendar by which initiatives aim to be ready to announce. For the international audience, it is Dubai's open door: a signal that the city's AI ambition is a shared venture rather than a closed government program, and a recruiting ground for the companies, researchers and builders the emirate wants alongside it. No entity is compelled to appear, yet the week fills, because presence there is where the city's AI conversation happens.





Instruments wielded, not owned

DCAI's restraint is most visible in how its instruments relate to the rest of this chapter. The Dubai AI Seal, examined next in these pages, is DCAI's trust signal to the private sector: a certification that ranks AI companies and opens the door to government partnership, run by the centre without regulating a single government entity. The AI Policy for government entities, covered later in this chapter, places compliance responsibility with each entity's own governance board and CAIO, with the Digital Dubai Authority engineering and governing the government-side architecture. And the AI+ Program, the citywide capability drive, is delivered by Digital Dubai with the Dubai Government Human Resources Department and DCAI. In each case the pattern repeats: DCAI mobilizes and steers at city level, while delivery and governance sit with the institutions built for them. The complementarity is deliberate. The centre supplies direction and momentum; the authorities supply the machinery.

The trade-off Dubai accepted

Coordination without control is a real trade-off, and Dubai made it deliberately. A centre that commands can guarantee uniformity; a centre that convenes must earn alignment, entity by entity, through the quality of its blueprint and the value of its forums. Dubai judged the price worth paying, because the alternative costs more. Commanded uniformity produces compliance without conviction, and AI adoption at city scale runs on conviction: on 27 entities that treat their AI portfolios as their own ambitions rather than someone else's instructions. The evidence of this report suggests the wager is paying off. The entities profiled in Chapter 3 built their systems themselves, on shared foundations, inside a common direction, which is precisely the outcome a coordinating institution exists to produce. Centralization and autonomy, the tension that recurs throughout this chapter, is resolved here at its source: centralize the direction, and leave the building to the builders.

From pivot to architecture

Dubai's State of AI Report 2025 described DCAI's launch as the moment the city moved from isolated projects to a coordinated, citywide AI effort, and credited the centre's foundational role in driving the AI agenda forward. One year on, that coordination has hardened into architecture. The chapters of that architecture follow directly from here: the Seal that recruits the market, the mandate that concentrates the mission, the policy and frameworks that set the rules, and the platforms, skills and compute that supply the means. Every one of them is easier to build, and to trust, because an institution at the centre of the city keeps them pointed in the same direction without holding any of them in its own hands.

> *"Dubai's answer to the coordination problem is not a super-ministry. It is a centre that convenes, accelerates and certifies, then trusts the city to build.*

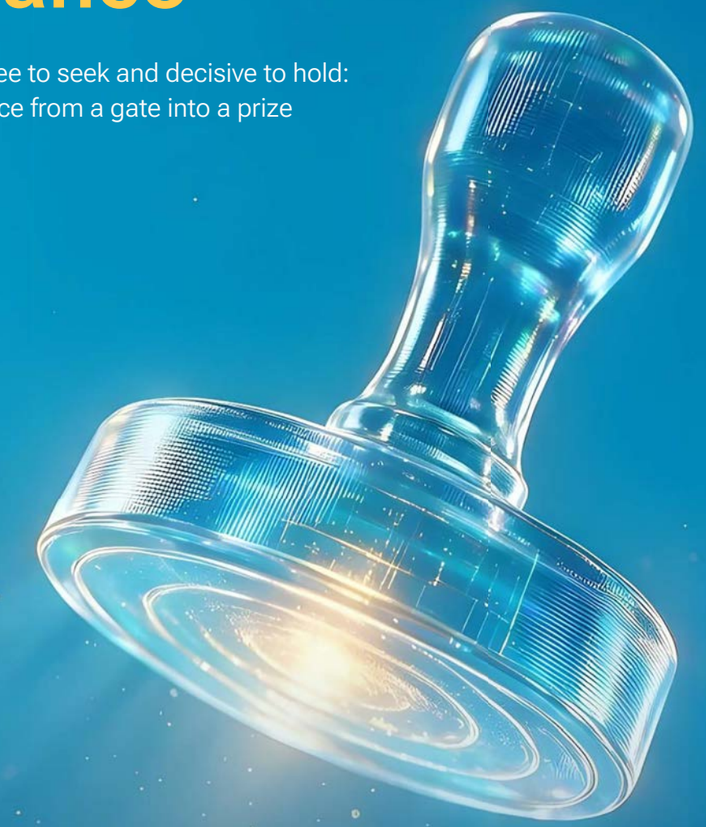
REFERENCE

[1] Dubai Future Foundation, "Dubai Centre for Artificial Intelligence," dubaifuture.ae, accessed Jul. 2026.
[2] Dubai Media Office, "Hamdan bin Mohammed launches Dubai Universal Blueprint for Artificial Intelligence," mediaoffice.ae, Apr. 2024.
[3] Dubai Media Office, "Dubai AI Week 2025 concludes," mediaoffice.ae, Apr. 2025.
[4] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

2 THE DUBAI AI SEAL:

Incentivized Governance

A certification that is free to seek and decisive to hold: Dubai turned governance from a gate into a prize



Most governance instruments work by subtraction: they define what cannot be done and penalize those who do it. The Dubai AI Seal works by addition. Launched under the directives of HH Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai, Deputy Prime Minister and Minister of Defence of the UAE, and Chairman of The Executive Council of Dubai, and aligned to the Dubai Universal Blueprint for Artificial Intelligence, the Seal is the Dubai Centre for Artificial Intelligence's trust instrument for the private market: a certification that identifies credible AI companies and rewards them with the most valuable signal the emirate can give, the standing to build with its government. Nothing about the Seal compels. Everything about it attracts. That inversion, governance as an incentive rather than a gate, is the idea this article examines, because it is the mechanism by which Dubai recruits a market instead of merely policing one.



Six tiers, E to S, rank companies by economic impact. The higher a company climbs, the more the city entrusts it to build.

How the Seal works

The design is deliberately simple. The Seal certifies AI companies, not government entities, across six tiers from E to S, ranked on the economic impact a company delivers in Dubai. Application is free, which removes the fee barrier that keeps smaller and younger firms outside comparable schemes elsewhere. Each certified company receives a verifiable serial number, so that a claim to hold the Seal can be checked rather than merely asserted, giving procurement teams, partners and customers a single trustworthy reference point. And because the tiers are graduated rather than binary, the Seal is not a pass-or-fail stamp but a ladder: a company certified at an entry tier has a visible route upward as its footprint and impact in the emirate grow. The choice of economic impact as the ranking dimension is itself telling. Dubai is not grading companies on the sophistication of their slide decks or the size of their global brand; it is grading them on what they contribute to the emirate. A young firm with real deployments in Dubai can outrank a famous name with none, which keeps the ladder honest and the incentive pointed where the city wants it: toward building here.

325
companies applied for certification

77
international offices represented

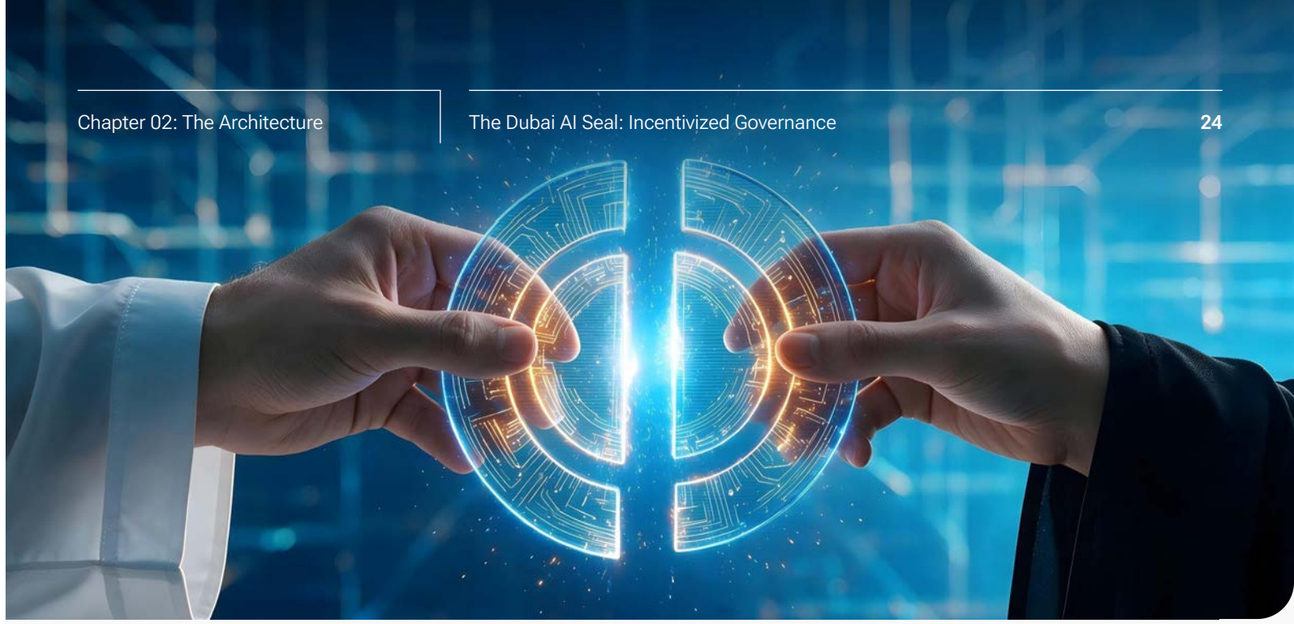
Not because the Seal is compulsory for them, but because government partnership runs through it.

The market answered

An incentive is only as good as the response it draws, and the response has been emphatic. By the count published in the Seal's first cycle, 325 companies applied for certification, together representing 77 international offices, a measure of how far beyond the emirate's borders the signal carried. At the summit of the first cohort, e& and IBM were among the first recipients of Tier S, the Seal's highest grade. The pairing is itself a message: one of the region's foremost technology groups and one of the world's longest-standing enterprise AI companies, certified side by side under the same Dubai standard. For the hundreds of applicants behind them, the top of the ladder is now visibly occupied, and visibly worth climbing. The geography of the applicant pool matters as much as its size. Presence across 77 international offices suggests the certification is read as durable well beyond the emirate. The applicant numbers demonstrate significant market interest in the certification; the data alone does not establish the causes of that interest.

The incentive engine

What makes the Seal more than a badge is the door it opens. Certification is required for companies partnering on AI projects with the UAE and Dubai governments. That single provision converts the Seal from a marketing asset into a business necessity for any firm that wants a share of a public-sector AI program of this scale. The economics are deliberate: the model shifts part of the compliance incentive to prospective suppliers, reducing reliance on purely enforcement-led mechanisms. A company that wants government work presents itself for certification voluntarily, documents itself thoroughly and maintains its standing carefully. Dubai gains a pre-vetted, tiered and continuously refreshed register of the AI companies serious about the emirate, and companies gain a credential whose value is underwritten by the government's own demand. There is a second-order effect as well. Because the Seal travels with the company rather than the contract, its value extends beyond



government work: a certification earned to win public-sector partnerships becomes a trust signal in private-sector negotiations too, in Dubai and in the 77 international offices the applicant pool represents. The emirate's standard, in other words, is beginning to price itself into markets Dubai does not regulate at all, which is what a successful trust instrument ultimately does.

Innovation and regulation, reconciled

Every jurisdiction confronting AI faces the same tension: regulate hard and innovation relocates; regulate lightly and trust erodes. The Seal is Dubai's answer to that trade-off, and it is a characteristically commercial one. Rather than writing rules that constrain what AI companies may build, the emirate defined what a trustworthy AI company looks like and attached the city's demand to that definition. Standards rise not because firms are compelled but because the most valuable customer in the market, the government itself, buys only from the certified. The result is a governance instrument that innovators run toward. It also fits DCAI's character, described in the previous article, of coordination without control: the centre does not tell the market what to build, it tells the market what the city trusts, and lets ambition do the rest. In the language of this chapter's operating model, the Seal is how Dubai mobilizes: the trust signal that recruits the market to build alongside the city, at the market's own initiative and the market's own cost of effort, with the city's standard as the price of entry.

The other half of trust

The Seal governs the market side of Dubai's AI compact; it deliberately leaves the government side to instruments of its own. What discipline looks like inside government, the mandate that concentrates every entity on a single platform, the policy that sets each entity's governance obligations, and the frameworks that will govern the agentic era, is the story the rest of this chapter tells. Read together, the two halves complete each other: a certified market building with a governed government. That is the compact this report calls scaling with trust, and the Seal is its opening move.

REFERENCE

[1] Dubai Centre for Artificial Intelligence, "Dubai AI Seal," dub.ai/en/ai-seal, accessed Jul. 2026.

[2] Dubai Future Foundation, "Dubai Centre for Artificial Intelligence," dubaifuture.ae, accessed Jul. 2026.

[3] Zawya, "Dubai AI Seal certification coverage," zawya.com, 2025.

[4] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

3 ONE PLATFORM, ONE YEAR:

The Mandate

On 1 April 2026, the Crown Prince gave every government service a single destination and Digital Dubai one year to coordinate the journey



Every individual and business service, every entity, one AI-powered platform, one year: a consolidation timetable of rare compression

Some directives adjust course. Others define an era. On 1 April 2026, HH Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai, Deputy Prime Minister and Minister of Defence of the UAE, and Chairman of The Executive Council of Dubai, issued one of the second kind: all Dubai government entities are to integrate every service for individuals and businesses into a single unified, AI-powered digital platform, within one year. Not a portal that links to entity websites, and not a shared design language stretched over separate systems, but one platform through which the government of one of the world's most dynamic cities serves everyone it serves. Digital Dubai was tasked with coordinating implementation across every agency. Measured against the scale of what must converge, the directive is a public-sector consolidation effort of rare scope, and the emirate handed the keys to Digital Dubai.

Why now: the technology finally fits the ambition

Dubai has pursued unified government before; what changed is the machinery available to deliver it. The platform is designed around generative AI and autonomous agents, and the difference is more than technical. Traditional integration forces every service into one rigid mould, which is why such programs historically stall. An AI-powered platform inverts the problem: generative interfaces let residents state what they need in their own words, while autonomous agents navigate the complexity behind the scenes, carrying a request across whatever entities it touches. The integration burden shifts from the citizen to the system. That is what makes a one-year horizon conceivable at all: the platform does not require every entity to become identical, it requires them to become reachable, by agents built to do the reaching. At the time of writing the platform is in build; the architecture described here is the intended design, and the instruments beneath it are the subject of the articles that follow.

The coordinator's task

Naming Digital Dubai coordinator concentrated accountability in the institution that has spent years assembling the prerequisites. Coordination here is an engineering assignment as much as an administrative one: a shared platform demands common standards for data, identity, security and, increasingly, agent behaviour, and it demands the means to apply them across dozens of entities at once, without flattening the autonomy that makes each entity effective in its own domain. It is also a test of trust in both directions. Entities must expose their services to a shared front door while retaining ownership of what happens behind it; the coordinator must set standards firm enough to guarantee coherence and light enough to leave room for the entity expertise the platform depends on. The balance between speed and safety, and between central coherence and entity autonomy, runs through every design decision the mandate now forces. One deadline, one coordinator, every entity: the directive turns a decade of digital-government groundwork into a single, dated deliverable.

The economics of one platform

The directive came with a projection attached. Official coverage of the announcement placed the platform's contribution at more than AED 10 billion in additional GDP within two years of full deployment, a return generated by faster transactions, lower friction for business, and the compounding value of services that anticipate needs rather than await applications. The platform is likewise positioned as a principal engine of the emirate's longer-term goal of an AED 100 billion digital economy. Both figures are projections rather than results, and this report prints them as such. What they signal is the frame in which Dubai holds the mandate: not an IT consolidation, but an economic strategy with a deadline.

Bringing the builders in

Alongside the directive, the Dubai-it initiative convenes the ecosystem that will help deliver it, bringing technology partners into the program rather than leaving the platform to government alone. A build of this scale and speed was never going to be delivered inside the government perimeter only, and the emirate did not pretend otherwise: the market certified through the Dubai AI Seal has, in this platform, its clearest invitation to build. The directive and the Seal thus meet in practice: one defines the work, the other defines who is trusted to share in it.

The mandate answered

A directive of this scale is only as credible as the machinery beneath it. The remainder of this chapter is that machinery: the policy that disciplines how entities adopt AI, the deployment playbook that sequences it, the decision-intelligence engine, the workforce program, the sovereign compute, the agents already at work in government's engine room, and the framework written to govern them. Read what follows as a single answer to a single question: how does a city deliver one platform in one year, without surrendering trust to speed?

AED 10B+ in additional GDP projected within two years of full deployment

AED 100B on the road to a digital economy of this scale



REFERENCE

[1] Dubai Media Office, "Hamdan bin Mohammed directs the integration of all government services into a single unified AI-powered digital platform," mediaoffice.ae, 1 Apr. 2026.

[2] Arabian Business, "Dubai to unify all government services on one AI-powered platform," Apr. 2026.

[3] Gulf News, "Dubai orders all government services onto a single AI platform within a year," Apr. 2026.

[4] Emirates 24|7, "Dubai-it initiative announced alongside unified platform directive," Apr. 2026.



4 The Rulebook That Made the City Faster

Dubai's AI Policy for Government Entities turned eleven ethical principles into binding operational mandates, and the first compliance cycle shows governance taking root at speed

> *Dubai designed governance as an accelerator: fourteen policies that tell every entity what good looks like, so none has to invent its own definition before it can move.*

The rulebook that arrived before the rush

In March 2025, a full year before the emirate's unified-platform mandate raised the stakes for every government entity, the Digital Dubai Authority issued the AI Policy for Government Entities: fourteen policies that convert Dubai's ethical commitments on artificial intelligence into binding operational mandates for the whole of government. The premise runs against the international grain. Where many administrations still treat AI governance as a brake to be applied once deployment is already under way, Dubai designed it as an accelerator: a common rulebook that tells every entity what good looks like, so that none has to invent its own definition of responsible AI before it can move. The policy applies to every Dubai Government entity, from departments and public corporations to the authorities supervising free zones, and it explicitly encourages the private sector to adopt the same comprehensive approach. It positions itself as three things at once: guidance for holistic adoption, a digital innovation mechanism for services and operations, and a means of institutionalizing AI across government, rather than a compliance checklist bolted on after the fact.

Fourteen policies, three layers

The fourteen policies are arranged in three layers. Four governance policies set the foundations: the eleven AI principles with which every government AI system should comply; the requirement that each entity establish an AI Governance Board; DDA's mandate to issue further AI regulation as common demand emerges; and continuous performance and impact assessment of AI systems in production. Seven implementation policies then guide the build itself: research and benchmarking, systems engineering, strategic alignment with entity and city strategy, shared services, prioritized use cases, collaborative partnerships, and rigorous testing across the full lifecycle. Three enabler policies address the human dimension: a culture of innovation and continuous learning, public outreach on AI achievements, and readiness assessment with a phased implementation roadmap. A dedicated annex extends the framework to generative AI, covering content attribution, accuracy safeguards, privacy in training data, and mechanisms of redress for anyone affected by generated content.

Eleven principles, one standard

At the centre of the framework sit the eleven principles: lawful, privacy preserving, human-centred, safe and secure, inclusive, fair, accountable, transparent, explainable, high-performance and robust, and sustainable. Each is defined in operational rather than aspirational language, and each reaches beyond the entity's own walls: any third party engaged on a government AI system, whether partner, supplier or contractor, must be equally bound to the principles through the agreements that engage them. The set mirrors the leading international references while remaining specific to Dubai's context, including the expectation that AI systems respect and reflect the cultural, linguistic and societal values of the UAE. DDA holds the regulatory pen, with the right to issue and modify the principles for AI systems across Dubai Government, while implementation and assurance rest with the entities themselves.

Accountability by design

The policy names its owners. In every entity, implementation responsibility rests with the Chief AI Officer, the same executive cadre that the Dubai Centre for Artificial Intelligence convened at city level. Oversight belongs to the entity's AI Governance Board, whose remit spans principles compliance, risk management, performance management with impact assessments conducted before and after deployment, stakeholder engagement, and AI skills development across the workforce. Entities may designate an existing committee or establish a new one: the policy mandates the function, not the organization chart. Compliance runs on the same logic of trust with verification. Entities self-assess against the principles, and DDA reserves the right to conduct selective compliance assessments. The regulatory stance is explicitly anticipatory and adaptive, with further policies, standards and guidelines to be issued as common demand from entities matures, so the regime can evolve as fast as the technology it governs.

Scope discipline

The policy is equally deliberate about what it does not govern. Data governance for AI systems remains the province of the regulations, standards and frameworks issued by the Dubai Data and Statistics Establishment, and AI security remains within the purview of the Dubai Electronic Security Center, including its Dubai AI Security Policy. Rather than duplicating either regime, the AI Policy preserves the separation of the three: data rules, security rules and AI rules compose cleanly, and no entity faces contradictory instructions from its regulators.



CALLOUT

Enforcement as an escalation ladder

The enforcement design deserves its own reading. Self-assessment is the default, keeping the compliance burden proportionate for entities at very different starting points. Selective assessment by DDA is the reserve power, held rather than routinely exercised, alongside the authority's right to stipulate compliance timeframes and methods. The policy was written to allow uneven maturity across the government, on the view that a single forced compliance posture would have slowed the leaders without helping the rest. Enforcement, in this design, is less a penalty regime than a synchronization mechanism.



The first evidence base

The 2025 self-assessment cycle, the first conducted under the policy, gives Dubai an aggregate picture of how quickly governance is taking root. Within roughly a year of issuance, 80 per cent of entities reported an active AI Governance Board, and 76.7 per cent had mapped their AI initiatives to entity and city strategy, the policy's central alignment test. More than half, 53.3 per cent, were already using AI shared services, and 66.7 per cent were engaged in AI research and benchmarking activity. Across the government, entities reported 30 AI partnerships and 44 AI applications tracked under the policy's assessment regime. This report publishes those figures in aggregate only: the self-assessment exists to synchronize the system, not to rank its members, and no entity's individual compliance position is disclosed.

2025 AI Policy self assessment: governance taking root

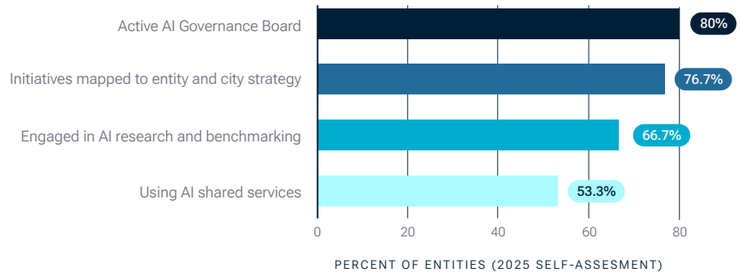


Figure 2.4: 2025 AI Policy self-assessment: governance taking root.

Speed and safety, reconciled

The classic objection to governing early is that rules slow builders down. Dubai's first compliance cycle suggests the opposite settlement of the speed-versus-safety trade-off. By the time the emirate's one-platform mandate landed in April 2026, four in five entities reported an active AI governance board, a shared vocabulary of principles, and a standing accountability chain from the Chief AI Officer upward. The rulebook did not slow the city; it synchronized it. What the policy leaves open, deliberately, is the question of method: how an entity should actually move an AI system from idea to production. That is the work of the deployment playbook on the pages that follow.

> *"Within roughly a year of issuance, 80 per cent of entities reported an active AI Governance Board. The rulebook did not slow the city; it synchronized it.*

REFERENCE

[1] Digital Dubai Authority, "AI Policy," internal document (Policy Version 1.0), Dubai, UAE, 2025.

[2] Digital Dubai Authority, "2025 AI Policy compliance self-assessment, aggregate results," internal document, Dubai, UAE, 2025.

[3] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

5

Four Quadrants, One Playbook

Publicly launched in April 2026 with more than 100 tools already live, Digital Dubai's AI Integration Matrix Framework gives every government organization a route from idea to production



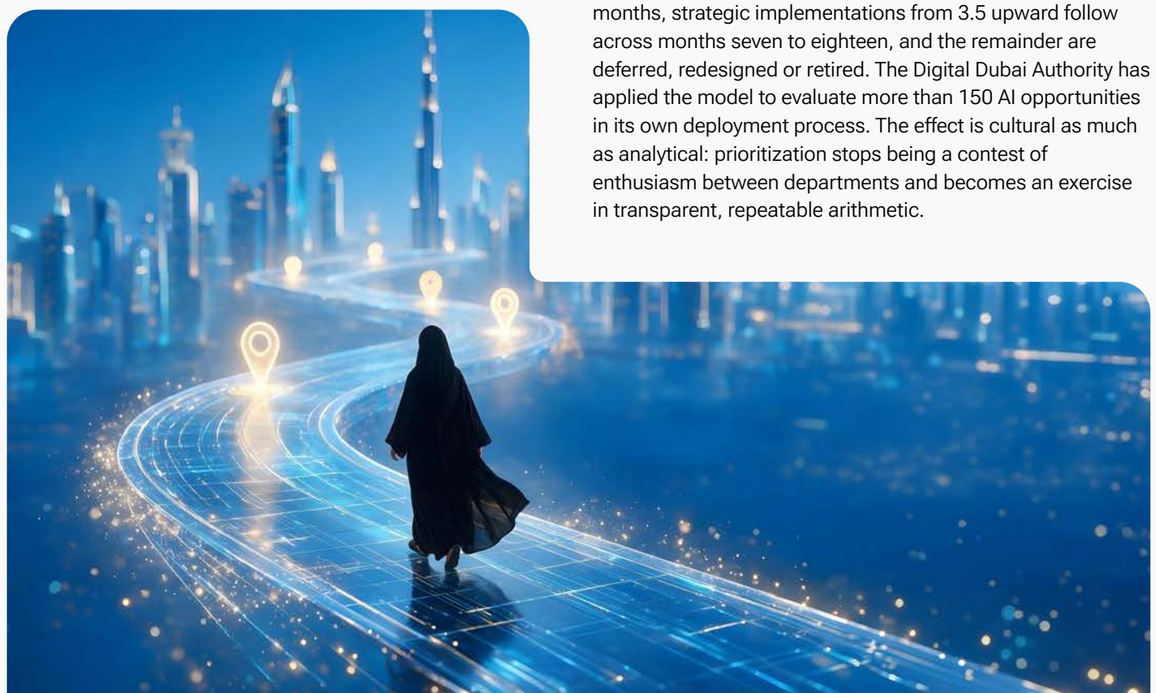
From internal doctrine to public playbook

On 28 April 2026, Digital Dubai publicly launched the AI Integration Matrix Framework, the structured model through which Dubai government organizations move artificial intelligence from idea to production. The launch, covered as the blueprint for embedding AI across Dubai's government, arrived with a proof point attached: more than 100 AI tools already live, and an explicit link to the emirate's shift toward a unified digital ecosystem. The framework had been developed at the Digital Dubai Authority in August 2025 and tested against the authority's own portfolio before it was offered to the city. It describes itself as a strategic framework rather than a formal policy document: where the AI Policy tells entities what responsible AI must be, the Matrix tells them how to build it, in what order, and with what governance around it. For the departments and authorities now working to a twelve-month deadline to unify services on a single AI-powered platform, that distinction is practical. The playbook is the fastest route from mandate to production.

Two questions, four quadrants

The model's power lies in its economy. Two questions sort every AI system a government could conceivably build. First, who does it serve: internal staff and back-office operations, or citizens and external stakeholders? Second, what does the technology do: act or answer? Systems that act are autonomous agents, executing tasks and decisions independently within defined rules. Systems that answer are retrieval-augmented generation (RAG) systems, combining advanced search with language generation to deliver precise, context-aware responses from existing knowledge. Crossing the two dimensions yields four strategic quadrants. Internal Agents automate government operations: routing documents, monitoring deadlines, processing invoices, matching talent to roles. Internal RAG systems put policies, procedures and institutional knowledge at every employee's fingertips. External Agents automate citizen service delivery end to end, from applications to transactions. External RAG systems answer citizens' questions about government services with accuracy and context. Each quadrant carries its own design principles, governance model and success metrics, and the framework insists that all four develop in balance: a government that automates only its citizen-facing layer while neglecting its own operations, or the reverse, is building an unbalanced portfolio.

> *"Two questions sort every AI system a government could build: who does it serve, and does it act or does it answer?*



Five steps from map to roadmap

Adoption follows a structured five-step process. Organizations first assess their current AI projects against the four quadrants, exposing gaps, overlaps and imbalances at a glance. They then evaluate each opportunity's strategic value, implementation complexity and risk profile, and prioritize a sequence that matches objectives, resources and stakeholder needs. A rollout roadmap is developed to address all four quadrants holistically rather than one quadrant at a time, and performance tracking closes the loop, monitoring progress and keeping initiatives coordinated as the portfolio grows. Alongside the process, the framework prescribes the change management that carries it: training, stakeholder engagement, internal communications and citizen support, on the principle that adoption is won in people before it is won in platforms.

Prioritization by arithmetic, not appetite

Classification alone does not settle what to build first. The framework's weighted prioritization model scores every candidate system on five criteria: impact potential, weighted at 30 per cent; implementation feasibility at 25 per cent; strategic alignment at 20 per cent; resource availability at 15 per cent; and risk at 10 per cent. Each criterion is scored from one (minimal potential or major barriers) to five (outstanding potential with minimal barriers), and the weighted sum produces a single comparable score for every opportunity in the portfolio. The scores then drive a phased roadmap: quick wins scoring 4.0 or above are sequenced into the first six months, strategic implementations from 3.5 upward follow across months seven to eighteen, and the remainder are deferred, redesigned or retired. The Digital Dubai Authority has applied the model to evaluate more than 150 AI opportunities in its own deployment process. The effect is cultural as much as analytical: prioritization stops being a contest of enthusiasm between departments and becomes an exercise in transparent, repeatable arithmetic.



PRIORITIZATION WEIGHTS

Impact 30%	Feasibility 25%	Strategic 20%	Resources 15%	Risk 10%

MATURITY MODEL

1 • Initial Awareness | 2 • Planned Development | 3 • Systematic Implementation | 4 • Optimized Operations | 5 • Predictive and Proactive

Source: AI Integration Matrix Framework whitepaper, Digital Dubai Authority, 2025 [1]

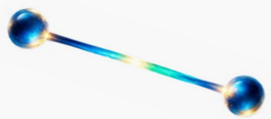
A maturity model for the whole journey

The framework pairs the matrix with a five-level maturity model that tells an organization where it stands and what to fix next. Level 1, Initial Awareness, describes isolated experiments and a reactive posture. Level 2, Planned Development, brings coordinated strategy, basic governance and first systems in multiple quadrants. Level 3, Systematic Implementation, marks comprehensive coverage across all four quadrants with mature governance processes. Level 4, Optimized Operations, adds continuous optimization and advanced analytics across an integrated estate. Level 5, Predictive and Proactive, is the destination: AI that anticipates needs and addresses challenges before they occur. The model is deliberately diagnostic rather than ceremonial, guiding targeted action on governance, talent, infrastructure and policy instead of awarding grades.

> *The framework guides more than 100 AI systems at the Digital Dubai Authority alone: the playbook was proven at home before it was offered to the city.*

The first customer was the author

The framework's credibility rests on the fact that DDA ran it on itself first, and the whitepaper's account of the starting point is candid: engineers wanted to build, leadership lacked a shared map, and initiatives risked fragmenting into silo solutions. The framework reversed that from the top down. Today it governs more than 100 AI tool already live across all four quadrants: 35 internal agents, 15 internal RAG systems, 40 external agents and 12 external RAG systems, each mapped to a defined role and strategic contribution. The portfolio runs through a documented governance stack: an AI Steering Committee of senior leaders that convenes monthly to approve enterprise initiatives and align them with national digital transformation goals; an AI Program Office acting as the central hub for standards, roadmap and vendor strategy; quadrant-specific working groups meeting every two weeks to carry practical insight from the ground upward; and an Ethics and Compliance Board that evaluates every implementation for ethical alignment, privacy safeguards and regulatory adherence. The whitepaper records the institutional effect in three registers. Management gained a precise map from which to set direction. Engineers gained clarity on exactly what types of systems to create, within one connected ecosystem rather than as silo solutions. And employees across the authority gained a clear line of sight from their own work to the citizen services it supports.



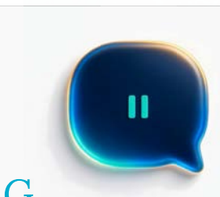
A2A

Agent-to-agent communication

Autonomous agents exchange data and tasks directly, so permit, verification, and payment systems can serve a citizen as one integrated service.

The frontier the matrix already anticipates

The framework is built for what comes next as well as what exists now. Its integration methodology insists that quadrants must not evolve in isolation, and it names the technology that will bind them: agent-to-agent (A2A) communication, in which autonomous agents exchange data and tasks directly, so that permit, verification and payment systems can serve a citizen as one integrated service. On the knowledge side, it points to conditional answer generation (CAG), a more disciplined evolution of RAG in which a system declines to answer rather than speculate when its information is insufficient, a safeguard tailored to government accountability. Both concepts prefigure the federated agentic architecture that closes this chapter.



CAG

Conditional answer generation

A disciplined evolution of RAG in which a system declines to answer rather than speculate when its information is insufficient.

The playbook's place in the architecture

Within Dubai's AI operating model, the Matrix sits squarely in the governance layer, and it completes that layer. The AI Policy defines the obligations; the Matrix supplies the method. Together they mean that a government entity facing the one-platform mandate does not begin with a blank page: it begins with four quadrants to map, a scoring model to sequence, a maturity level to diagnose and a governance stack to copy. Few governments anywhere publish their deployment playbook. Dubai launched its own, with the evidence of more than 100 live tools attached, and invited every government organization to use it.

REFERENCE
[1] Digital Dubai Authority, "AI Integration Matrix Framework for government organizations," whitepaper, Version 1.0, Dubai, UAE, 2025.
[2] Gulf Today, coverage of the public launch of the AI Integration Matrix Framework by Digital Dubai, 28 April 2026.
[3] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

6 DDA by the Numbers

One authority's measurable footprint at the centre of Dubai's government AI architecture: every figure dated and sourced.



42%
average monthly growth
in the running total of tokens processed on the government LLM platform, llmapi.govai.ae
SOURCE: AI Infrastructure Empowerment Platform, since Sep 2025 to Oct 2026 average of the month-on-month increases

100+
live AI systems governed
by the AI Integration Matrix Framework
SOURCE: Framework public launch, Apr 2026

AED 125M+
verified savings
delivered by CODE.AI
SOURCE: EC Resolution 44/2024 monitoring, 2025

50,000
government employees
entering AI training by 2030
SOURCE: AI+ Program, Apr 2026

140+
interactive dashboards
in live use
SOURCE: CODE.AI, 2025

40
AI Champions
leading adoption inside Digital Dubai
SOURCE: AI Connect, Mar 2026

33
AI agents
built or in development on the GRP platform
SOURCE: Shared Services Applications Department, Nov 2026

80%
of entities
with an active AI governance board
SOURCE: 2025 AI Policy self-assessment, aggregate

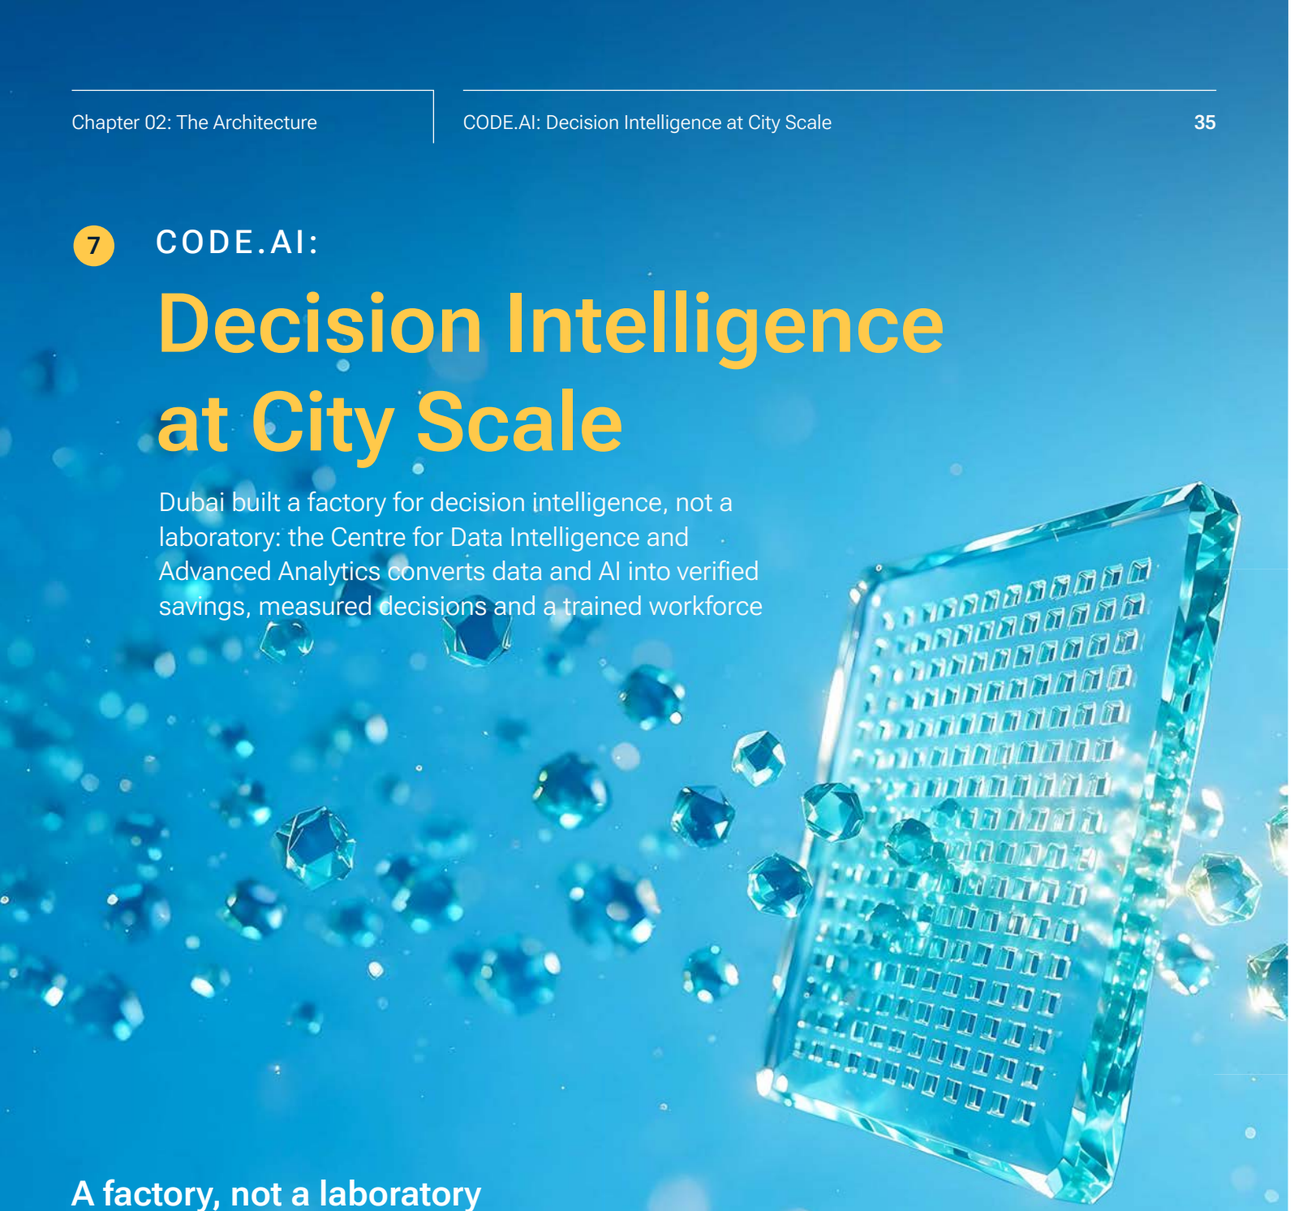
All figures verified at source; aggregates published without entity attribution.

REFERENCE
[1] Dubai Digital Authority, "AI Infrastructure Empowerment Platform: running total of tokens processed by month, September 2025 to October 2026," internal records, Dubai, UAE, 2026.
[2] Gulf Today, coverage of the public launch of the AI Integration Matrix Framework by Digital Dubai, 28 April 2026.
[3] Digital Dubai Authority, "CODE.AI Executive Council Resolution 44/2024 monitoring form," internal document, Dubai, UAE, 2025.
[4] Digital Dubai, "AI+ Program launch release," Dubai, UAE, 15 April 2026.
[5] Digital Dubai Authority, "AI Connect digest," internal document, Dubai, UAE, 2026.
[6] Dubai Digital Authority, Shared Services Applications Department, "Smart Employee services: AI agents," internal document (Arabic), Dubai, UAE, 2026.
[7] Digital Dubai Authority, "2025 AI Policy compliance self-assessment, aggregate results," internal document, Dubai, UAE, 2025.

7 CODE.AI:

Decision Intelligence at City Scale

Dubai built a factory for decision intelligence, not a laboratory: the Centre for Data Intelligence and Advanced Analytics converts data and AI into verified savings, measured decisions and a trained workforce



A factory, not a laboratory

Most governments respond to the AI moment by building a lab. Dubai built a factory. CODE.AI, the Centre for Data Intelligence and Advanced Analytics, was established as a network initiative under Executive Council Resolution 44/2024, accountable to the Digital Dubai Authority and aligned to the Dubai Data and Statistics Strategy. Its mandate is not to demonstrate what AI might one day do for government decision making but to industrialize it: to embed AI and advanced analytics in the decisions the city actually takes, and to report the value in numbers an auditor would accept.

The design was researched before it was built. CODE.AI's founding team consulted 16 Dubai Government entities and benchmarked 16 analytics functions worldwide, from Singapore's whole-of-government data science division to the United Kingdom's data science campus, London's cross-entity coordination model and Estonia's data strategy. The consultation found no shortage of ambition across the city's entities; it found a coordination gap. Collaboration was concentrated on data sharing when it needed to extend to algorithms, compute and scarce human expertise. CODE.AI was designed to close that gap as a network, with entities as members rather than clients.

CODE.AI'S EVIDENCE BASE

140+
interactive dashboards in operation

6,000+
statistical indicators tracked

19
entities at the highest data-maturity level

Four pillars, one discipline

The initiative operates through four pillars. Decision Intelligence embeds AI and advanced analytics across government decision making. Growth and Governance provides the services that keep adoption wide and responsible, including ethical AI and model assurance. Active Innovation supplies thought leadership and experimentation, from a government-wide sandbox to the Dubai Methods Unit. Workforce Empowerment upskills the government workforce on AI and data. Four pillars, one discipline: decisions supported by evidence, value measured after the fact.

Savings against a target

That discipline is most visible in how CODE.AI reports its own impact. The initiative has delivered more than AED 125 million in verified savings against a target of more than AED 500 million by 2028: AED 85 million from onboarding entities to the Dubai Data Platform, AED 29 million from the Dubai AI Platform, and AED 11 million from interactive dashboard development. The framing matters as much as the figure. These are verified savings set against a dated target, not an aspiration. The same reporting covers the wider evidence base: more than 140 interactive dashboards in operation, more than 6,000 statistical indicators, and the highest data-maturity levels reached in 19 government entities. The Dubai Dashboard now gives government leadership central access to strategic and performance indicators and the analytical insight behind them, so the same evidence discipline that produced the savings also carries the city's day-to-day decisions.

CODE.AI verified savings: AED 125M+ towards the AED 500M target by 2028



Figure 2.7: CODE.AI verified savings: components, 2025 total and the AED 500M target by 2028.



Verified savings of AED 125 million and counting, against a target of more than AED 500 million by 2028

CODE.AI measures itself the way it asks the city to decide.

The PoC Factory: a conversion pipeline

CODE.AI's best-known service is the PoC Factory, and the second word is the important one. It is a factory, not a showroom: a shared methodology that takes a use case from a requesting entity, gathers requirements, allocates the right engineering resources, builds and tests, and then applies a value-measurement gate that weighs benefits against effort before a stop-or-go decision. Use cases that pass the gate move toward production; those that do not are stopped early and cheaply, and the learning is banked. The measure of the Factory is conversion into production deployments, not the volume of experiments it hosts. Early use cases carried through the pipeline include a population prediction model built on household water and electricity consumption data and a generative AI assistant

supporting policy development, alongside synthetic labour-force survey data that lets analysts work on sensitive datasets while preserving privacy, security and confidentiality.

Around the Factory sits a widening portfolio: a government-wide AI and advanced analytics sandbox, advisory and enablement services delivered to six government entities, the city's first pan-government strategic project built on data intelligence, and a research partnership with the Rochester Institute of Technology that has delivered two use cases.



> *"Six leadership sprints have taken 49 executives from 36 government entities through the same discipline: frame the priority, interrogate the evidence, commit.*

Leaders first, then the workforce

CODE.AI treats decision makers as its primary users. Six leadership sprints have taken 49 executives from 36 government entities through a compressed discipline: frame the strategic priority, interrogate the evidence, evaluate the options and commit to a course. In parallel, more than 600 government staff have been trained on AI and advanced analytics toward a target of 7,000, a thread that now runs into the city-wide AI+ Program covered in the next article.

600+
government staff trained on AI and advanced analytics

7,000
CODE.AI target for government staff trained on AI and advanced analytics

From centre of excellence to network

The operating model has evolved deliberately: from a centre of excellence, through a consulting posture, toward the target factory-and-network state in which pod teams pair CODE.AI engineers and data scientists with entity product owners and subject-matter experts. That evolution is Dubai's answer to the centralization-versus-autonomy trade-off in decision intelligence. The scarce capabilities, data science, engineering and assurance, are pooled where they can reach every entity; the decisions, and the ownership of use cases, stay with the entities themselves. It is the same pattern this chapter traces in compute (article 2.9) and in agents (articles 2.10 and 2.11): centralize the means, distribute the choices. CODE.AI is where the pattern was first proven against a savings target, and it is the measurement engine that keeps the rest of the architecture accountable.

REFERENCE

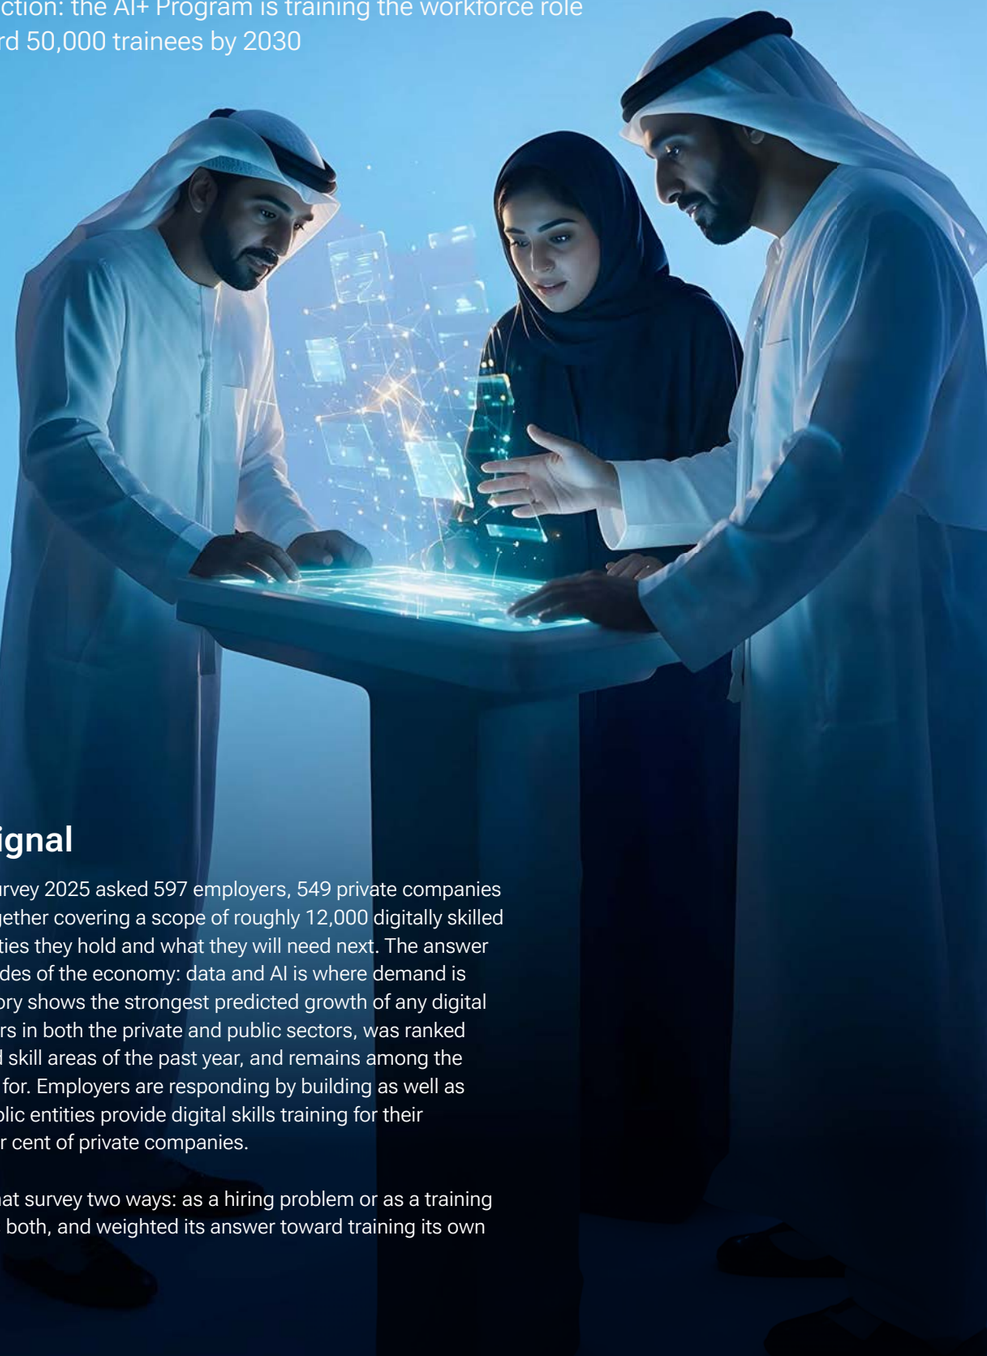
[1] Digital Dubai Authority, "CODE.AI: A network initiative to leverage AI and advanced data analytics for Dubai," internal document, Dubai, UAE, 2024.

[2] Digital Dubai Authority, "Executive Council Resolution 44/2024 monitoring form: Centre for Data Intelligence and Advanced Analytics," internal document, Dubai, UAE, 2026.

[3] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

8 AI Workforce & Upskilling Architecture

Dubai treats AI as a core skill for every government employee, not a specialist function: the AI+ Program is training the workforce role by role, toward 50,000 trainees by 2030



The demand signal

The Dubai Digital Skills Survey 2025 asked 597 employers, 549 private companies and 48 public entities, together covering a scope of roughly 12,000 digitally skilled employees, what capabilities they hold and what they will need next. The answer was consistent on both sides of the economy: data and AI is where demand is concentrating. The category shows the strongest predicted growth of any digital skill over the next two years in both the private and public sectors, was ranked among the most recruited skill areas of the past year, and remains among the harder capabilities to hire for. Employers are responding by building as well as buying: 96 per cent of public entities provide digital skills training for their employees, against 71 per cent of private companies.

A government can read that survey two ways: as a hiring problem or as a training mandate. Dubai read it as both, and weighted its answer toward training its own people at scale.

AI+: one program, five tracks

On 15 April 2026, Digital Dubai launched the AI Workforce Transformation Program (AI+) in collaboration with the Dubai Government Human Resources Department (DGHR) and the Dubai Centre for Artificial Intelligence, under the umbrella of the Dubai Future Foundation. It is a multi-year capability program with a single positioning decision at its core, drawn directly from Dubai's AI Policy: AI is a core skill for every employee, not a specialist function housed in an innovation unit.

The program reaches the workforce role by role through five tracks. The Leadership track runs expert-led sessions and roundtables on the opportunities and limits of AI in government, with case studies from leading countries. The Chief AI Officers track focuses on implementing Dubai's AI policy and framework, the technical enablers provided by Digital Dubai, and practical case studies from government entities, delivered through workshops and communities of practice. The Directors track equips departmental managers to identify and scale high-impact use cases, apply AI-driven decision making and practise responsible AI governance. The Product and Service Owners track carries participants from ideation to applicable use cases and through to impact measurement and value realization. The General Employees track builds working skills for every desk: prompt engineering, automating tasks with AI agents, and responsible use, delivered through webinars, in-person training and virtual courses.

> *From 240 trainees in 2025 to more than 1,000 employees across 78 entities in 2026, toward 50,000 by 2030: capability built on a curve, not a course.*

Arabic, free, and ending in working agents

The Directors track shows the program's operational character most clearly. It is delivered in Arabic and free of charge, running function-by-function cohorts through 2026: human resources, marketing, procurement, customer happiness, strategy and policy, finance, and legal. The training day is only the opening move of a structured 90-day journey. Within 21 days, each director identifies and prioritizes AI use cases inside the department's own work. Each director then nominates a champion from the team, and the champion is trained to build AI agents for those use cases, with the director guiding selection and clearing the path. At the 90-day mark, the director measures the change in team productivity and decision speed and feeds the results into the program's impact assessment. The unit of completion is not a certificate; it is a working agent and a measured result.

The arc to 50,000

The trajectory is deliberate. In 2025, 240 trainees completed AI literacy training in eight batches delivered with Microsoft: the proof that the government workforce would engage. In 2026, AI+ scales to more than 1,000 employees across the 78 entities in the program's scope, spanning all five tracks, with managers and the wider employee base phased in through 2027. By 2030, the target is 50,000 trained Dubai Government employees. Each stage changes the question being asked: 2025 asked whether employees would learn, 2026 asks whether departments can convert learning into use cases, and the years to 2030 ask whether an entire government can make AI fluency as unremarkable as email.







"In Dubai, the government employee is no longer merely a user of technology, but a key partner in designing smart, proactive services."

H.E. Hamad Obaid Al Mansoori
Director General, Digital Dubai

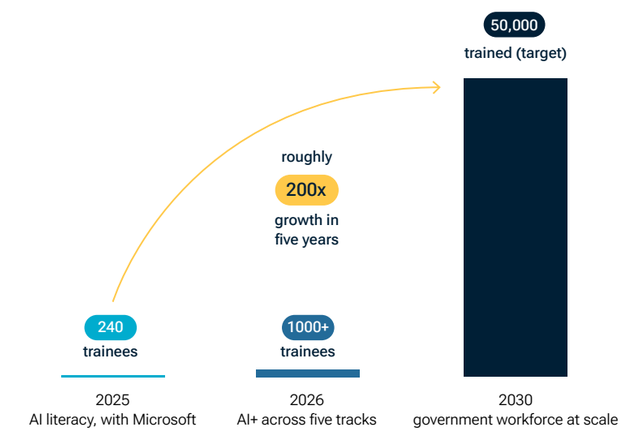


Figure 2.8: The AI+ arc: from 240 trainees to 50,000 by 2030.

At the launch, H.E. Abdulla Ali bin Zayed Al Falasi, Director General of DGHR, described investing in people as "the foundation for shaping the future", while H.E. Khalfan Belhoul, CEO of the Dubai Future Foundation, framed the program as a milestone in turning advanced technologies into practical, high-impact use cases across government entities. Three institutions, one curriculum: the delivery authority, the human resources regulator and the city's AI coordinating centre.

Workforce as architecture

AI+ completes a load-bearing wall in the operating model this chapter describes. The sovereign compute layer (article 2.9) supplies the tokens; the shared GRP platform (article 2.10) supplies agents; the policy and the Integration Matrix (articles 2.4 and 2.5) supply the rules and the playbook. None of it moves without people who can specify a use case, prompt a model and judge an output. DGHR's own transformation appears in its entity spotlight in Chapter 3; here it stands as co-architect of the government-wide skills layer. The wager is specific and testable: by the time the unified platform mandate falls due, the scarcest input to Dubai's AI architecture will not be compute or models but trained judgement, and Dubai will have manufactured it at scale.

REFERENCE

[1] Digital Dubai, "Digital Dubai launches 'AI+' to build an AI-ready government workforce," digitaldubai.ae, 15 Apr. 2026.
[2] Digital Dubai, "Dubai Digital Skills Survey Results 2025," Dubai, UAE, 2025.
[3] Digital Dubai Authority, "AI+ Program: HR orientation," internal document, Dubai, UAE, 2026.
[4] Dubai Government Human Resources Department and Digital Dubai, "AI+ Program: Directors Track brochure (Arabic)," Dubai, UAE, 2026.
[5] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

9

INFRABOOST & THE AI INFRASTRUCTURE EMPOWERMENT PLATFORM:

Sovereign Compute for Government

Dubai pooled the government's demand for AI compute into one sovereign inference layer: shared infrastructure, free model choice, and a running token total growing 42% a month.



A coordination problem first

Every government entity that wants to put large language models to work faces the same shopping list: GPUs, model hosting, a security review, API plumbing, and an operations team to keep it all running. When dozens of entities work through that list separately, the city pays for the same learning many times over, capacity sits stranded in silos, and the security perimeter multiplies with every deployment. Dubai's reading of this is the founding argument of Infraboost, its strategic program for aggregating government AI infrastructure: GPU access at city scale is a coordination problem first and a procurement problem second. Solve the coordination, pool the city's demand into shared infrastructure with a common operating discipline, and procurement follows on better terms, once, instead of many times badly.

In October 2025, HH Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai, Deputy Prime Minister and Minister of Defence of the UAE, and Chairman of The Executive Council of Dubai, approved the AI Infrastructure Empowerment Platform, giving the shared-infrastructure approach its official mandate. What had begun as an engineering program became city policy: government AI workloads would run on sovereign, shared, centrally operated

infrastructure. Six months later, the directive of 1 April 2026 to unify all government services on a single AI-powered platform raised the stakes further. A platform that will lean on generative AI and autonomous agents for an entire city's services needs an inference layer that is already provisioned, already secured and already governed. Dubai had spent the previous year building exactly that.

The shared inference layer

The working core is llmapi.govai.ae, the Dubai Government's shared inference platform, live since September 2025. It exposes AI models through OpenAI-compatible APIs running on high-performance GPU clusters, which means any entity system built against the industry-standard interface can connect without rework, and any application written for the platform can switch models without rebuilding. Alongside this large language model service, the platform offers GPU as a Service: entities that need raw compute can apply for direct, time-bound access to the GPUs themselves. The numbers since launch describe adoption, not intent: 31 government entities onboarded, and a running total of tokens processed that has grown by an average of 42 per cent a month.

The model catalogue is the platform's quiet statement of sovereignty. All 19 deployed models are self-hosted on government infrastructure. According to DDA's architecture specification, model inference runs on the government's own clusters and prompt and response content is not retained by the platform. The catalogue spans eight model families: text, omni, vision, embeddings, rerankers, document OCR, text-to-speech and decision. The catalogue grows as open models improve; the interface entities build against does not change. Self-hosting open models is the hinge of the design: it delivers frontier-class capability while keeping the weights, the logs and the hardware inside the government's own perimeter.

> *"31 government entities, 19 self-hosted models, a running token total growing %42 a month: one shared sovereign inference layer, live since September 2025*

Privacy by architecture

The platform's privacy posture is structural rather than contractual. User prompts pass through the API gateway and are discarded after processing; nothing is saved. Model outputs are transient: responses are never logged, cached or persisted. What the platform keeps is metadata only: user identity, prompt length, token count, timestamps, throughput and latency, for audit and capacity planning. Prompt content, model responses and conversation history are never tracked. For government workloads that touch personal or sensitive data, this converts a policy assurance into an architectural fact: there is no stored corpus of government prompts to secure, because none exists. Governance monitoring runs on the metadata alone, which is sufficient to answer the questions oversight actually asks: who is using the platform, how much, and how well is it performing.

> *"Prompts are never stored and responses are never logged. The platform keeps metadata for audit and capacity planning, never content.*

Shared-inference economics

The economic argument for the platform is an argument about utilization. AI inference demand is spiky: a statistics team's month-end crunch, a customer-service peak, an agent workflow that runs hot for an hour and sleeps for six. An entity that provisions its own GPU estate must size it for its own worst hour and watch it idle the rest of the time; multiply that across dozens of entities and the city is warehousing idle silicon. A single shared platform aggregates those spikes: one entity's quiet hours absorb another's peak, and every additional tenant improves the utilization of the same clusters. This report deliberately prints no savings figure for that difference. A credible number needs a denominator, the counterfactual cost of fragmented entity-by-entity deployment, and that baseline is still being established. The growth of the running token total stands as the interim evidence: demand at this pace, materializing within months of launch and returning daily, is the behaviour of a service offered at no cost to the entities that use it.

Running shared infrastructure concentrates risk as well as value, and the program names its own risks plainly: balancing capacity so GPU resources are neither starved nor stranded; securing shared APIs against unauthorized access, data leakage and misuse; and operating at scale, with orchestration, monitoring and high availability across the estate. Naming them is part of the governance discipline, and each maps to a standing control: utilization monitoring for capacity, key-based access and gateway controls for security, and the metadata-only logging regime that gives operators observability without ever holding content.

Centralize compute, free the model choice

The deepest design decision in Infraboost is where centralization stops. Compute, security, privacy and operations are pooled, because these are the layers where fragmentation is pure waste. Model choice is not pooled. Entities choose among the catalogue's models and categories according to the use case in front of them: an OCR-heavy licensing workflow, an Arabic-language citizen assistant and a text-to-SQL analytics tool have different needs, and the platform serves all three without imposing a house model. This is the centralization-versus-autonomy trade-off resolved the same way the chapter has already seen it resolved in decision intelligence (article 2.7) and in skills (article 2.8): centralize the means, distribute the choices. The centre runs the racks; the entities keep the judgement.

What this makes possible

The platform's significance runs past the applications it serves today. Agentic systems, in which AI agents plan, call tools and complete multi-step work with limited supervision, consume tokens at a rate conversational use never approached. A government that metered every entity's agent workloads through separate commercial endpoints would find the economics prohibitive and the governance opaque. A government that runs agents against a shared sovereign inference layer, with metadata-level observability and key-based access already in place, can afford to let agents multiply, and can watch them do it. That is the ground on which the next two articles stand: the GRP agents already modernizing the government's back office (article 2.10), and the federated framework being prepared to govern the agentic era (article 2.11). Infraboost makes Dubai's federated agentic architecture economically possible. The framework that follows makes it safe.

42%
average monthly growth
SOURCE: in the running total of tokens processed, llmapi.govai.ae, September 2025 to October 2026

31
government entities onboarded
SOURCE: llmapi.govai.ae, September 2025 to October 2026

19
self-hosted models
SOURCE: llmapi.govai.ae, September 2025 to October 2026

EIGHT MODEL FAMILIES

SOURCE: llmapi.govai.ae, October 2026

- **Text**
- **Omni**
- **Vision**
- **Embeddings**
- **Rerankers**
- **Document OCR**
- **Text-to-speech**
- **Decision**

TWO WAYS TO USE IT

LLM as a Service
Ready-to-use models through the API with no infrastructure to manage and standard configurations.

GPU as a Service
Direct and time-bound access to GPU resources, subject to approval.

REFERENCE

[1] Digital Dubai Authority, "AI Infrastructure, Security, and Integration Standards," internal document, Dubai, UAE, 2026.
[2] Digital Dubai Authority, "Infraboost: scalable AI infrastructure project deep dive," internal document, Dubai, UAE, 2025.
[3] Dubai Media Office, coverage of the approval of the AI Infrastructure Empowerment Platform, mediaoffice.ae, Oct. 2025.
[4] Dubai Media Office, "Hamdan bin Mohammed directs the integration of all government services into a single unified AI-powered digital platform," mediaoffice.ae, 1 Apr. 2026.
[5] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

10

GRP AI AGENTS:

Modernizing Government Back Office Operations

On the shared Government Resource Planning platform, an agent built once can serve every entity, and new agents are being built to enter invoices, prepare payments and reconcile accounts



The invisible engine room

Citizens never see it, and that is rather the point. Behind every government service sits a back office: invoices to verify, ledgers to reconcile, suppliers to onboard, shifts to schedule, candidates to screen. In Dubai, that machinery does not live in dozens of separate systems. It lives on the shared Government Resource Planning (GRP) platform, the common finance, procurement and human resources backbone that Digital Dubai develops and manages for more than 80 government entities [1]. Digital Dubai's portfolio of AI services for GRP now counts 33 agents across those three functions, built or in development, designed to work inside the systems employees already use [2].

The strategic consequence is easy to state and hard to overstate. Because the platform is shared, an agent built once can be offered to every entity on it. There are no parallel procurements, no duplicated integration projects, and no maturity gap opening up between the largest department and the smallest authority. The reconciliation service, for example, is switched on with a support request, not a procurement [3]. One engine room powers the whole fleet.

Thirty-three agents, three functions

In finance, agents match receipts to invoices, check invoices against system data and tax rules, reconcile bank statements and supplier balances, and review VAT treatment before a return is filed. In procurement, they guide suppliers through registration, answer supplier questions in Arabic and English, draft requests for proposal and contracts, and score bids for tender committees. The largest group serves people: a digital assistant answers employees' questions on their records, laws and policies and helps them submit requests, while other agents explain payslips, build shift rosters, rank CVs, run first interviews and turn skill gaps into learning plans. An assistant that drafts employees' objectives from their manager's scorecard is in testing [4]. In September 2026, Dubai Police and Digital Dubai launched an AI-assisted shift management system in the Smart Employee app, which Digital Dubai plans to offer to entities across Dubai [1]. In the catalogue's own descriptions, these agents draft, score, match and recommend; a person or committee with the authority makes the final decision.

Three agents that complete the work

The three finance agents in development in 2026 go further: they complete the work inside GRP [5], [6]. The invoice entry agent reads a scanned invoice, checks the supplier, purchase order, VAT and duplicates against GRP's records, and creates the invoices that pass. The payment agent takes a request in plain language, retrieves the eligible invoices, explains why each one is prioritized, excluded or flagged, and passes the approved batch to the bank through a direct host-to-host link. The reconciliation agent builds on the live service, which compares sub-ledger and general ledger balances and recommends fixes. It matches individual transactions across GRP and bank records, clears high-confidence matches where governance rules allow and routes the rest with a reason code.

> *"One shared platform means an agent built once can serve finance, procurement and human resources teams in every entity on it.*

One control model

All three agents share one control model [5]. Every value or match the AI infers carries a confidence score, and low-confidence items are highlighted for review. AI judgement is kept separate from GRP's fixed business rules, so an invoice that fails a rule goes to a person, however confident the model is. Exceptions arrive with the reason and a suggested correction, and an audit trail links each transaction to its source document, the agent's reading and every human change. People approve every payment run and review every material difference. In the GRP team's words, "AI handles interpretation and prioritization; GRP remains the system of record." That is the platform's answer to the tension between automation and human oversight.

Targets first, results next

In a year, 167,552 supplier invoices were keyed in by hand and 211,138 supplier payments were prepared across the platform, work the GRP team puts at more than 136,000 staff hours [7]. The targets are set: 60 to 80 per cent of invoices created with no manual step, 70 to 90 per cent of payment preparation automated and 80 to 95 per cent of reconciliation matches made automatically [5]. These are targets, not results. The GRP team scales each agent only when accuracy, control coverage and payback thresholds are met [5].

When agents act

The GRP agents settle, in miniature, the tension that runs through this chapter: centralization versus autonomy. The platform centralizes what is identical everywhere, the mechanics of reconciling, verifying, matching, drafting and scheduling, while each entity keeps what is properly its own: judgement. That division is easy to hold while agents recommend. An agent that creates an invoice or prepares a payment run is no longer advising an officer; it is working inside the government's accounts. That changes the work of the teams who run the back office, whom AI+ is training function by function (article 2.8). It also raises the question the next article answers: when agents act rather than advise, and dozens become hundreds working across entity boundaries, who writes the rules? Dubai has already answered, in writing.



AI handles interpretation and prioritization; GRP remains the system of record.



IN DEVELOPMENT: AGENTS THAT COMPLETE THE WORK

Targets: invoice entry, 60 to 80% of invoices with no manual step · payment entry, 70 to 90% of preparation automated · reconciliation, 80 to 95% of matches made automatically

REFERENCE

[1] Gulf News, "Dubai Police deploy AI to make workforce planning smarter," gulfnews.com, 10 Sep. 2026.
[2] Dubai Digital Authority, Shared Services Applications Department, "Smart Employee services: AI agents," internal document (Arabic), Dubai, UAE, 2026.
[3] Dubai Digital Authority, "Financial Reconciliation service: launch announcement and user manual," internal documents, Dubai, UAE, 2026.
[4] Dubai Digital Authority, "SmartEmployee SmartPath: employee objective creation with AI assistant," internal presentation, Dubai, UAE, 2026.
[5] Dubai Digital Authority, "Agentic finance operations: AP invoice entry, AP payment entry and financial reconciliation," executive presentation, Dubai, UAE, Nov. 2026.
[6] Dubai Digital Authority, "AI use cases: AP Invoice Entry, AP Payment Entry and Financial Reconciliation agents," internal document, Dubai, UAE, 2026.
[7] Dubai Digital Authority, "GRP AI use cases 2026: ROI calculations," internal workbook, Dubai, UAE, 2026.

11 THE FEDERATED AGENTIC AI FRAMEWORK:

Governing the Agentic Era

With agents already in production, Dubai wrote the rulebook for the agentic era before scaling it: a governance framework for a city of cooperating AI agents, and practical guidelines for every single one



Proof first, then policy

The previous pages described AI agents at work on a shared government platform, some now being built to complete transactions. That is precisely the moment, as governments worldwide are discovering, when governance becomes urgent: after autonomy has been demonstrated, but before it is everywhere. In June 2026, Digital Dubai Authority completed two instruments built for exactly this window. The Federated Agentic AI Governance Framework sets the rules for a city in which autonomous agents, operated by different public and private entities, cooperate at scale. The Guidelines for AI Agents give every Dubai Government entity a practical playbook for each individual agent it deploys. Together they make Dubai one of the first cities anywhere to publish a dedicated rulebook for the agentic era ahead of mass adoption rather than in reaction to it.

Both documents are explicitly anticipatory. They establish principles, guardrails and governance mechanisms in advance of market maturity, oriented toward plausible future use cases as much as present ones, so that entities can adapt with agility as conditions evolve. The intent is constructive rather than restrictive: autonomy balanced with oversight, innovation with control, so that AI agents act with integrity and measurable alignment to public trust.

A federation, not a monolith

The framework's core idea is federation. Dubai's agentic future will not be one giant AI system but many autonomous agents cooperating across organizational boundaries while each retains local governance, decision-making capability and accountability. The framework's ambition is that this multitude behaves as one trusted system: shared registries, synchronized policies and cross-entity decision logging keep the federation coherent while innovation stays local. Scope discipline is deliberate. Security regulation for AI remains with the Dubai Electronic Security Center (DESC), and sector-specific rules remain with sector regulators, preserving the clean division of institutional labour described earlier in this chapter.

Nine principles, in plain words

The framework rests on nine governance principles. Subsidiarity: decisions are taken at the lowest competent level, with entities governing their own transactions and escalating only when required. Proportional data with privacy preservation: an agent receives only the minimum information needed for its approved purpose, and sensitive data stays where it is generated. Transparency, traceability and accountability: every decision in the federation must be explainable, linked to policy and provenance, and attributable to a responsible entity, with evidence generated at the moment of execution rather than reconstructed later. Interoperability by design: shared schemas and message standards let agents from different entities transact predictably as one ecosystem. Proportionality of governance controls: low-risk actions are governed lightly to preserve efficiency, while rights-affecting actions face stronger enforcement and human review. Unified accountability across the federation: decision-making is decentralized, but roles and responsibilities at entity, cross-entity and city level are explicit, and responsibility is shared without ever being diluted. A single governance charter with local autonomy: one set of rules for all participants, with diversity of implementation beneath it. Registry-anchored authority: policy execution is validated against a Federated Registry, so only approved policies can run and revoked ones are blocked. And human judgment as final arbiter: autonomous action escalates to people wherever citizen rights, safety, conflicting policies or ethical questions arise.

> *"Dubai's answer to the agentic era is anticipatory: principles, registries and guardrails established before widespread adoption, not after the first incident.*

The machinery behind the principles



Principles need machinery. DDA will act as Dubai's lead federation entity: it certifies entities and their agents into the federation, may revoke participation where assurance falls below acceptable thresholds, and manages the Dubai-wide Federated Registry, the authoritative record of participating entities, agents, policies, schemas and performance benchmarks. A cross-entity Federated Agentic AI Taskforce, with one designated member from each participating entity, anchors oversight, escalation and shared semantics inside every organization. The framework then commits DDA to five companion instruments that translate principle into engineering: an architecture blueprint for the federation's layered, policy-aware design; an Interoperability Framework so agents communicate in a secure, verifiable and policy-aligned manner; a Knowledge Management Framework, the shared semantic layer that keeps every agent grounded in a common, authoritative body of city knowledge; a Security and Trust Framework, developed in close collaboration with DESC; and a Performance and Quality Framework covering service levels, telemetry and drift detection across the federation.

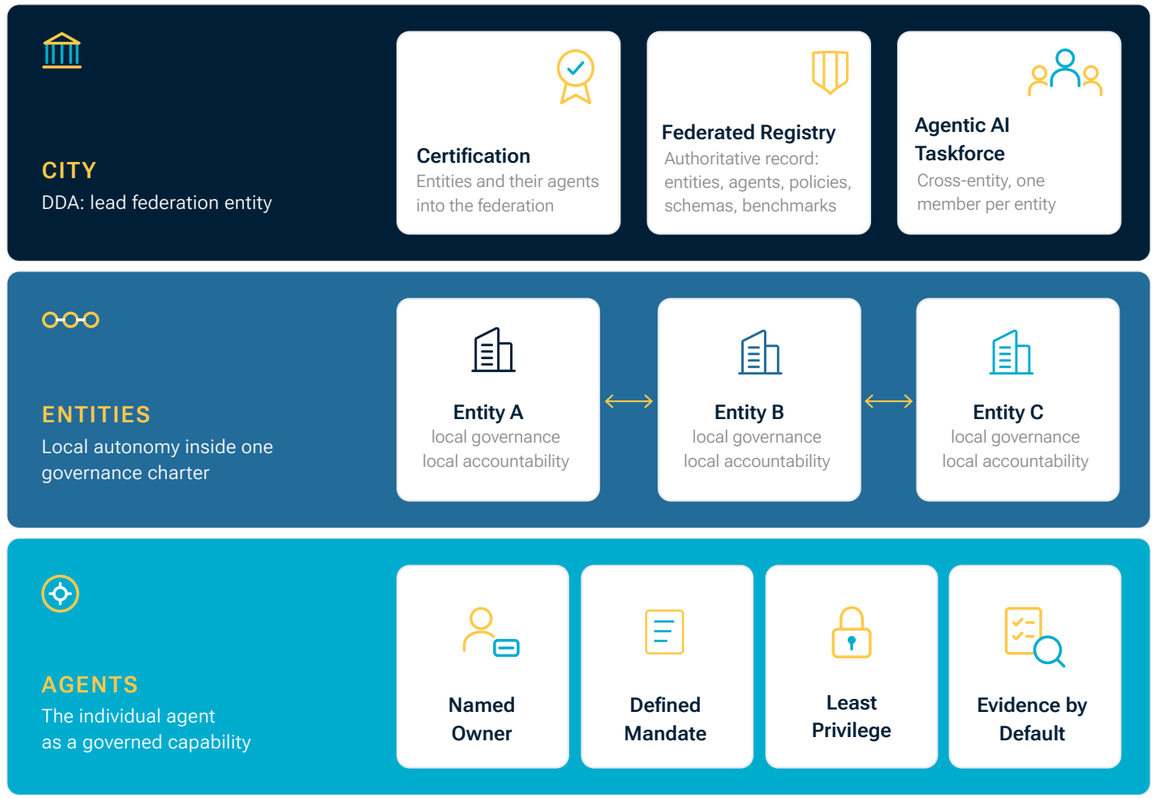


Figure 2.11: The federated governance architecture: city, entities, agents.

NINE GOVERNANCE PRINCIPLES

Subsidiarity
Decisions are taken at the lowest competent level

Proportional data with privacy preservation
Only the minimum information needed for the approved purpose

Transparency, traceability and accountability
Each decision can be explained, traced and attributed

Interoperability by design
Agents from different entities transact as one ecosystem

Proportionality of governance controls
Lighter for low-risk actions, stronger where rights are affected

Unified accountability across the federation
Responsibility is shared without ever being diluted

Single governance charter with local autonomy
One set of rules for all participants

Registry-anchored authority
Only approved policies can run and revoked ones are blocked

Human judgment as final arbiter
People decide on rights, safety, policy conflicts and ethics

These nine principles govern the federation · the seven guiding principles of the Guidelines for AI Agents govern the individual agent

COMPANION FRAMEWORKS

- **Architecture blueprint**
- **Interoperability**
- **Knowledge management**
- **Security and trust with DESC**
- **Performance and quality**

Security regulation stays with DESC · sector rules stay with sector regulators

Rules for the single agent

Where the framework governs the federation, the Guidelines for AI Agents govern the individual agent inside a government entity. Their starting position reframes the technology: an AI agent is not merely software but a delegated operational capability, and delegation demands explicit accountability, a named owner, a defined mandate and clear boundaries. Seven guiding principles follow. Lawful authority: agents act only within the entity's legal mandate and approved purpose. Human accountability: identifiable officials remain responsible for an agent's design, operation and consequences. Meaningful human control: people can understand, approve, override and pause consequential actions. Rights protection: due process, accessibility and contestability are designed in from the outset. Least privilege: an agent receives only the data, tools and time windows its approved task requires. Evidence by default: material decisions and tool calls are logged in auditable form. Clarity and remedy: people affected by an agent know it, can reach a human, and can contest and correct outcomes.

The Guidelines add three instruments of proportionality. A five-tier risk classification runs from Tier 0, prohibited uses such as final decisions on citizen rights without meaningful human review, through low, medium and high risk, to Tier 4, critical agents with high autonomy touching payments, identity or city infrastructure, which require senior approval, continuous monitoring, mandatory human override and emergency shutdown. Four human oversight models let entities match supervision to stakes: human-in-the-loop approval before each material action, human-on-the-loop live supervision with intervention rights, human-over-the-loop periodic review of samples and metrics, and human-out-of-the-loop operation reserved for minimal-risk, tightly constrained workflows. And a nine-stage lifecycle, from use-case identification and risk classification through design, data and tool readiness, development, testing, controlled rollout and monitored operation to formal decommissioning, treats agents as managed capabilities rather than one-time deployments.

Many agents, many entities, one trusted system: local autonomy inside a single governance charter, with human judgment as the final arbiter.



A rulebook with room to grow

The framework anticipates its own extension. Sector regulators, such as the Dubai Health Authority in healthcare, the Knowledge and Human Development Authority in education and the Roads and Transport Authority in mobility, may in time choose to serve as lead federation entities for their own domains, overseeing sectoral federations of public and private agents under the same logic. That prospect points directly at the sector clusters of Chapter 3, where those regulators' production systems already run. It also completes this chapter's loop. The Architecture opened with governance and closes with it: Dubai scales what it has first made trustworthy. Scaling with trust is not a slogan attached to the operating model; it is the operating model.

REFERENCE
[1] Digital Dubai Authority, "Federated Agentic AI Governance Framework," internal document (v1.0), Dubai, UAE, 2026.
[2] Digital Dubai Authority, "Guidelines for AI Agents," internal document (v1.0), Dubai, UAE, 2026.
[3] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

12 AI CONNECT:

Digital Dubai's Internal Community of Practice

Before asking a city to adopt AI, Digital Dubai practises on itself: 40 AI Champions across 28 departments, meeting weekly to move use cases from challenge to implementation

Practising on itself first

Before an authority asks a city to adopt AI, it should be able to show what happened when it asked itself. AI Connect is Digital Dubai's internal community of practice, spanning the Digital Dubai Authority and its establishments (DDA, DDGE and DDSE): 40 AI Champions drawn from 28 departments, each serving as the liaison between their department's employees and the strategy and performance function on AI use cases. The community meets weekly, and the sessions have one job: moving use cases from challenge to implementation.

The cycle is deliberately operational. Champions are nominated by their departments; challenges are identified in the daily work; AI-driven solutions are proposed against them; the pipeline is prioritized and assessed; and implementation proceeds through collaboration and engagement across departments. Sessions cover both practical implementations, where a department has already expedited a process, reduced the resources an activity requires or introduced a new practice, and forward proposals: use cases that do not yet exist but would remove a known problem if built. The community states its reasons for investing the time plainly: increasing efficiency, eliminating duplicate cases, solving problems, sharing knowledge, and enabling collaboration across the organization. The weekly rhythm is the discipline that makes those reasons real. A use case cannot drift for a quarter, because someone will ask about it next week.



From internal habit to public products

The practice pays in two currencies. The first is efficiency inside the organization: duplicate efforts surface early and are merged, a solution built in one department is rolled out to all of them, and knowledge moves at the speed of a weekly meeting rather than an annual review. The second is the pipeline into Dubai's public-facing portfolio. Products the city knows by name, GovGPT, DubaiAI, SHAMMA and ClassifAI, pass through this same internal practice, where champions exercise capabilities, surface gaps and plan roll-outs before government entities and the public come to depend on them. A capability proven at a champion's desk earns its way outward, first across departments, then into the products the city uses. The organization is its own first user, and its own toughest one.

The close of the enable layer

AI Connect is deliberately modest in scope, and that is the point. It is an internal practice of Digital Dubai, not a government-wide network; the city-scale structures are the ones this chapter has already described. But it closes the ENABLE layer at human scale. The same organization that runs the city's sovereign inference platform, trains the government workforce and writes the deployment playbook also spends part of every week asking 40 of its own people what AI changed in their departments since the last session. That habit is the smallest structure in this chapter and among the most telling. An architecture is credible when its architect lives in it. Dubai's does.

REFERENCE

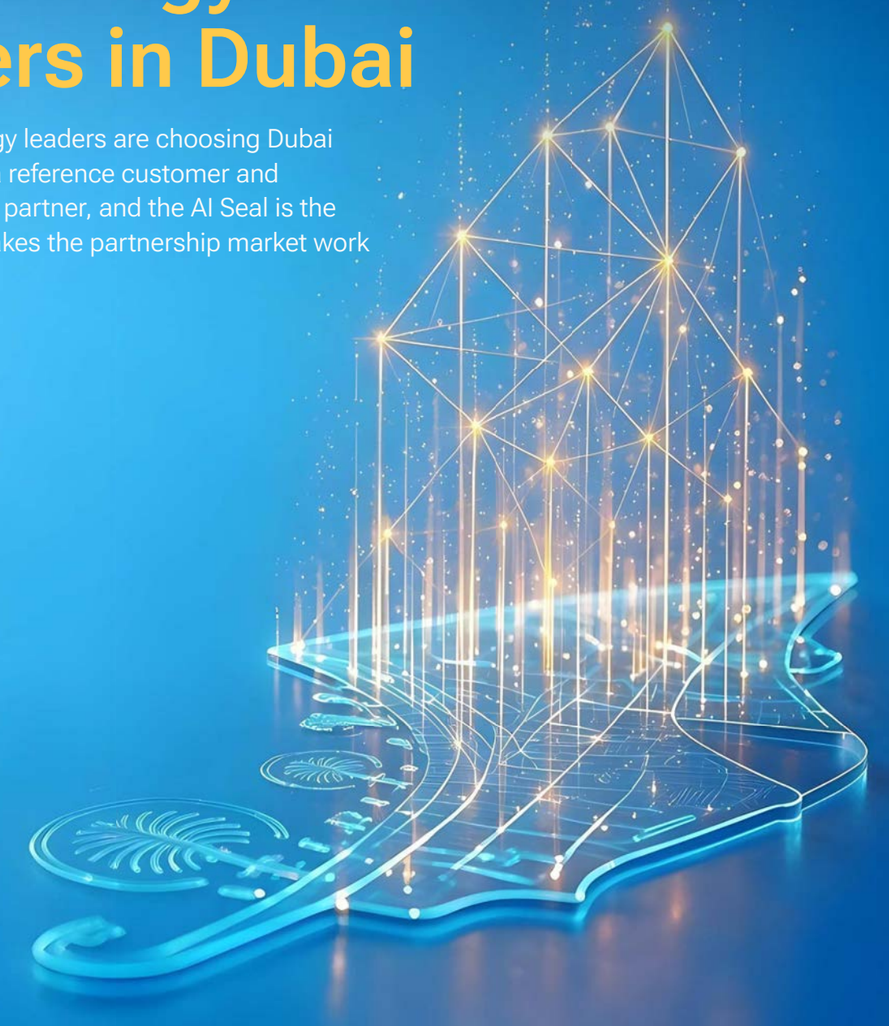
[1] Digital Dubai Authority, "The AI Connect: intelligent collaboration from within," internal document, Dubai, UAE, 2026.

[2] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

13 PRIVATE SECTOR SPOTLIGHT:

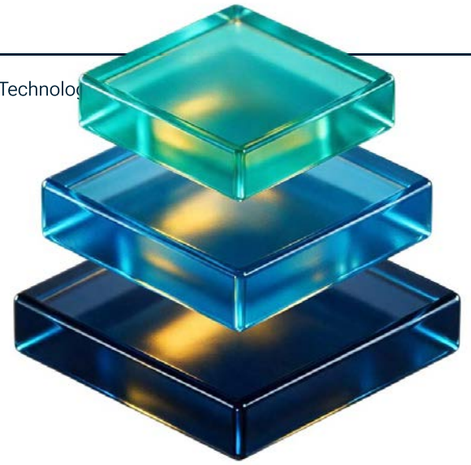
Technology Players in Dubai

Global technology leaders are choosing Dubai government as a reference customer and co-development partner, and the AI Seal is the gateway that makes the partnership market work



A partnership market, mapped

Read across the entity spotlights of Chapter 3 and a pattern emerges: the world's technology leaders are not selling to Dubai government at arm's length. They are building with it. The 2026 entity submissions name 37 distinct private-sector AI partnerships across ten of the report's twenty featured entities, spanning global platform providers, sector specialists, UAE national champions and homegrown firms. For these companies, Dubai government is not simply a customer; it is a reference deployment in one of the world's most visible digital administrations. What follows is a map of that ecosystem rather than an endorsement of any part of it, and the map has four distinct layers.



The platform layer

Microsoft appears more often than any other name: an enterprise AI foundation and Copilot deployment at DEWA, a Dynamics 365 and Power Platform rollout for citizen services alongside consolidated AI-driven security operations with Dubai Municipality, and work with the Community Development Authority. OpenAI's collaboration with DEWA has placed official utility services in the ChatGPT Apps Directory, and Google's NotebookLM powers DEWA's AI-generated public podcasts. SAP is co-developing a path toward autonomous ERP with DEWA, including agents tailored for the utility industry, while GitHub's Copilot underpins AI-assisted engineering there. IBM, one of the first two Tier S recipients of the Dubai AI Seal alongside e&, delivers its watsonx platform and data lakehouse at GDRFAD, where SAS provides advanced analytics and HPE the high-performance AI infrastructure. SAS also works with the Roads and Transport Authority on AI-driven traffic signal optimization, and Capgemini partners with the Department of Finance on AI-powered invoice recognition and touchless accounts payable.

The specialists

A second layer of partnerships brings deep sector expertise. In health, Epic embeds predictive risk models and generative AI in Dubai Health's electronic health records. In energy, DEWA draws on Bidgely for smart-meter intelligence, Solargis and Solcast for solar resource assessment and forecasting, and Fugro for AI-driven geospatial monitoring of utility infrastructure. In mobility, Swiftly powers predictive real-time passenger information for RTA journeys, delivered through the journey planners passengers already use, and Pony.ai brings its Level 4 autonomous driving platform to the emirate's roads. Emerging agentic tooling firms such as Genspark and n8n round out DEWA's orchestration stack. The pattern across all of them is the same: domain depth applied to a government platform that is ready to receive it.

National champions and the homegrown layer

UAE-built AI is entering city workflows on equal terms. Space42, the Abu Dhabi geospatial and satellite intelligence company, and AIQ, developer of the agentic ENERGYai platform for the energy sector, both feature in DEWA's partnership portfolio, evidence that the national AI champions scale at home first. Beneath them sits a productive layer of regional and local firms: Cyberspace, Diverge, Brixbey and Neol are building the Dubai Government Human Resources Department's AI portfolio, from Talent GPT candidate shortlisting to legislative drafting assistance; Talents Arena co-created the Dubai Fit AI Companion with Dubai Sports Council; Inffy and Techvista work alongside Microsoft at the Community Development Authority.

> *"Thirty-seven named technology partnerships across ten featured entities, and one certification gateway: the AI Seal turns market access into a governance incentive.*

The gateway: why they come

What converts this interest into an orderly market is the Dubai AI Seal, introduced earlier in this chapter. A company seeking to build with Dubai government does not negotiate its trustworthiness deal by deal; it certifies once, under the Dubai Centre for Artificial Intelligence's program, and carries a verifiable credential across the ecosystem. Certification across six tiers, E to S, is free to apply for and required for companies partnering on government AI projects; 325 companies have applied to date. The effect is a single trust gateway at the front door of a government AI program of this scale. For Dubai, the Seal converts governance into magnetism. For the technology industry, the message is equally clear: the route into Dubai's deployment market runs through demonstrated trustworthiness, and the queue at that gate is growing.

REFERENCE

[1] Dubai Government entity submissions, "State of AI Report 2026 entity data returns (private-sector partnership fields)," Dubai, UAE, 2026.

[2] Dubai Centre for Artificial Intelligence, "Dubai AI Seal," dub.ai/en/ai-seal, Dubai, UAE, 2025.

[3] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

14 PRIVATE SECTOR SPOTLIGHT:

Investment, Funding, Capital Flows

Capital follows clarity: Dubai's governance stack is doing for AI investment what its trade infrastructure once did for goods, removing uncertainty before money has to price it



The de-risking argument

Investment capital does not fear ambition; it fears ambiguity. The questions that stall AI investment decisions elsewhere are rarely about the technology: they are about the rules. Who certifies a vendor? Which framework governs an autonomous agent? Where does government-grade compute come from, and on what terms? What happens when a deployment touches citizen data? Every page of this chapter has been, in effect, an answer to one of those questions, and that is precisely what makes Dubai's institutional architecture an investment story. A published AI Policy, a launched Integration Matrix, a certification Seal, a sovereign shared-inference platform and a prepared rulebook for the agentic era each remove a variable that an investor would otherwise have to price as risk. Policy certainty is not a soft asset. It compounds.

Each element of the stack answers a different investor question. The AI Seal tells a company exactly how to qualify for government partnership, in six published tiers, with a free application. The Integration Matrix tells a vendor how a government entity will prioritize, score and govern the deployment it is proposing. The shared inference platform tells a product team what sovereign, government-grade AI capacity looks like before a contract is signed. And the agentic governance framework tells the market what tomorrow's autonomous products will be permitted to do, before those products exist anywhere. Regulatory foresight of that kind is normally the scarcest commodity in an emerging technology market. In Dubai it is published.

The demand signal

The market Dubai has created is a deployment market with a deadline. The directive of 1 April 2026, unifying all government services on a single AI-powered digital platform within one year, carries official projections of more than AED 10 billion in additional GDP within two years of full deployment, and sits inside the emirate's longer-term goal of a digital economy contributing AED 100 billion annually. Both figures are projections from official announcements, but the behaviour they describe is already observable. At GITEX Global 2025, Digital Dubai ran the Dubai Government pavilion with more than 50 participating entities and an explicit agentic AI focus: less an exhibition than an open procurement horizon. And the Dubai AI Seal has drawn 325 corporate applicants representing 77 international offices, each willing to submit to certification because a government partnership in Dubai is worth qualifying for. Companies do not queue at a gate unless something valuable sits behind it.

> *"A projected AED 10 billion in additional GDP within two years, an AED 100 billion digital economy goal, and a rulebook that tells investors exactly how the city works.*

37
named technology partnerships committed to the emirate

10
government entities involved

325 companies queued for AI Seal certification: the clearest demand signal a deployment market can send.



From venture bets to deployment capital

The composition of the opportunity matters as much as its size. Markets without institutional clarity attract early-stage speculation: small bets on possibilities. Markets with governed, funded, deadline-driven demand attract deployment-phase capital: production systems, co-development programs, multi-year platform commitments. Dubai's architecture pushes the opportunity decisively toward the second category. The 37 named technology partnerships across ten government entities described in the previous article are, from a capital perspective, evidence that global firms are committing delivery capacity, regional presence and co-innovation investment to the emirate, the corporate equivalent of capital expenditure rather than option-buying. The Dubai-it initiative, convening the technology community around the unified platform effort, extends the same invitation to the wider market.

The architecture as prospectus

For an investor or a corporate strategist weighing where to place resources in a crowded global AI market, then, this chapter can be read as a prospectus in institutional form. The mandate defines the demand and its timetable. The policy and the matrix define the operating rules. The Seal defines market access. The shared infrastructure defines the cost base. The agentic framework defines the permissible frontier. Few AI markets anywhere offer that sequence in writing, and none of it depends on any single deal closing. Capital, patient and otherwise, tends to notice when a city removes its excuses one by one. Dubai has been removing them methodically, in public, one framework at a time, and the next article shows what has already been built on the ground it cleared.

REFERENCE
[1] Dubai Media Office, announcement of the unified AI-powered government services platform directive, Dubai, UAE, 1 April 2026.
[2] Dubai Centre for Artificial Intelligence, "Dubai AI Seal," dub.ai/en/ai-seal, Dubai, UAE, 2025.
[3] Digital Dubai, Dubai Government pavilion at GITEX Global, Dubai, UAE, October 2025.
[4] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

15 PRIVATE SECTOR SPOTLIGHT:

Projects & Built Cases

Co-built systems already in production across the emirate show what the architecture delivers when a global partner plugs into a governed platform



Built, shipped, running

Strategies are announced; systems are shipped. Dubai's claim to be a deployment city rather than a strategy city rests on the second category, and the evidence is specific. Consider five systems, each co-built by a government entity with a named private partner, each in production today, and each drawn from a different sector of the city's operations: security, health, utilities, mobility and border services.

At Dubai Municipality, an agentic AI-driven Security Operations Centre built with Microsoft has cut cyber response from manual processes measured in hours to automated action measured in seconds, a performance leap rarely achieved even in mature, world-class security operations.

At Dubai Health, Epic's predictive clinical risk models and generative AI assistants run inside Salama, its Epic-based electronic health record, providing early insights that support timely care and helping clinicians streamline documentation, while ambient clinical documentation turns the patient-physician conversation into notes, giving physicians more time with patients.

> *"From AI-driven cyber defence to clinical risk models at the bedside: the proof of Dubai's architecture is what has already shipped.*

At DEWA, a collaboration with Google made the utility the first government entity globally to turn its press releases into AI-generated, podcast-style audio briefings through NotebookLM, and a parallel collaboration with OpenAI placed official DEWA services in the ChatGPT Apps Directory, giving customers secure, conversational access to the utility where they already are.

At the Roads and Transport Authority, work with SAS on the Marrakesh Corridor produced an AI optimization engine that generates tailored traffic signal timing plans for peak-hour conditions, paired with forecasting dashboards that fuse signal, navigation and sensor data for the emirate's traffic engineers.

And at GDRFAD, IBM's watsonx platform and data lakehouse underpin the analytics estate that powers proactive residency services, turning one of the world's busiest border and residency authorities into a predictive operation.

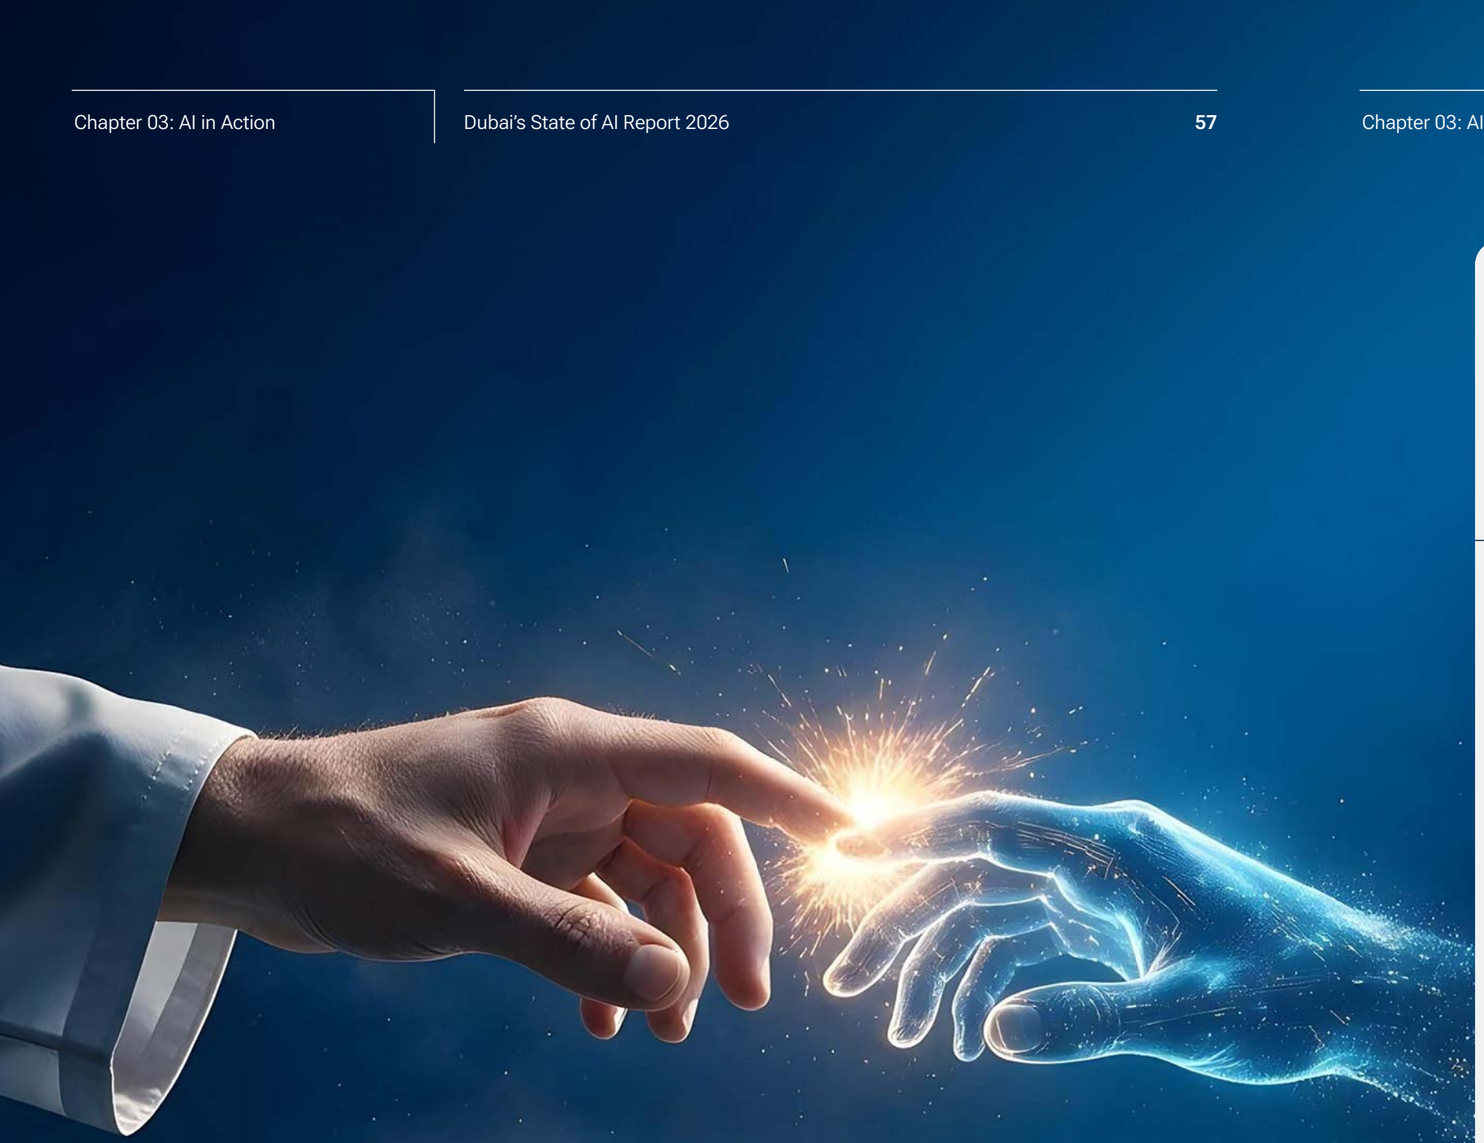
What the cases have in common

None of these systems is a demonstration, and none of them stands alone. Each pairs a government entity that owns the problem with a private partner that brings specialized capability, working on platforms governed by the frameworks this chapter has described and under the trust instruments it established. The partner gains a reference deployment in a global city; the entity gains capability it did not have to build alone; the city gains proof that its architecture converts partnership into production, on the terms its frameworks set. That exchange, repeated across sectors, is what the MOBILIZE layer of Dubai's AI operating model exists to produce.

From architecture to evidence

This closes the chapter, and the aperture now widens. Everything described here, the coordination, the mandate, the policies, the platforms, the skills, the agents, the rulebook and the partnerships, exists to produce results inside the institutions that run the city. The next chapter is that architecture in action, twenty featured entities deep: more than ninety production systems across seven sectors, told by the entities themselves. The rulebook was this chapter. The evidence is the next.

REFERENCE
[1] Dubai Government entity submissions, "State of AI Report 2026 entity data returns," Dubai, UAE, 2026.
[2] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.



AI in Action

The architecture described, the proof delivered: seven sectors, 20 entities, and the production systems now serving Dubai every day.

> *Chapter 2 described the architecture. This chapter is the proof: production systems, measured outcomes, and the residents and businesses they serve.*

Every claim in this report ultimately rests on one question: does it work in production? Chapter 2 set out the architecture, the governance frameworks, the data foundations, and the institutional machinery through which Dubai scales artificial intelligence with trust. Chapter 3 answers the question. Across the pages that follow, 20 government entities open their operations to show what that architecture delivers when it meets the daily business of a city.

Production, not promise

The rule of inclusion is deliberate and strict: everything featured here is live. These are systems in production, handling real transactions, serving real residents, and reporting measured outcomes, from approval times cut by weeks to hundreds of thousands of service hours returned to higher-value work. Where an entity is building foundations rather than running systems, the chapter says so plainly, because disciplined sequencing is itself a form of maturity.

How to read this chapter

The 20 entities are organized into seven sector clusters: government services, mobility and urban development, energy and utilities, health, security and safety, economic growth, and knowledge and human development. Each cluster opens with a leadership voice, a senior figure from within the sector reflecting on what the past three years have taught and where the next three lead. The entity spotlights for that sector then follow, each built on the entity's own submission: what is live, what it has changed, and where it goes next. Read in sequence, the clusters trace a single arc: an administration that has moved artificial intelligence from the edge of government to its core. The proof begins overleaf, with the entities that run the business of government itself.

1

GOVERNMENT SERVICES:

The Core Gets Intelligent



A leadership reflection from H.E. Abdulrahman Saleh Al Saleh, Director General of the Department of Finance, opens the cluster where government runs itself on AI

If artificial intelligence is to be judged anywhere, it should be judged here: in the machinery of government itself. The seven entities in this cluster manage the emirate's public finances, its government workforce, its borders and residency, its prosecution and judicial coordination, its endowments, and the housing of its citizens. These are the functions residents cannot opt out of, which is precisely why they are where trust must be earned first.

A reflection from the Department of Finance

For H.E. Abdulrahman Saleh Al Saleh, Director General of the Department of Finance, the period from 2023 to 2025 was a defining chapter: the department's evolution from early adoption to a structured, visionary roadmap. "We view AI not merely as a tool," he writes, "but as a foundational pillar for supporting the Emirate's leading economic position and ensuring long-term financial sustainability." The ambition is global leadership in sustainable financial innovation, with AI as a catalyst for more efficient public spending, equitable resource allocation, and a higher quality of life for the community.

The strategy reaches beyond back-office automation into the fabric of financial policy: predictive fiscal modelling, more accurate revenue forecasting, and evidence-based evaluation of legislative and policy impacts, all supported by robust data reuse and analytics that give leaders real-time insight. The journey, he reflects, has reinforced a central insight: important technology is only as effective as the governance framework supporting it. A Responsible AI approach, grounded in preventing bias, protecting privacy, and maintaining human accountability, proved essential to building the public trust that large-scale adoption requires.



Scaling demanded institutional transformation too: integrating advanced algorithms with legacy infrastructure, competing for scarce specialised talent, and holding data quality to the standard on which every model depends. Looking to 2026 and beyond, his commitment is to convert those challenges into opportunities through sustained development of digital infrastructure and human capital, keeping Dubai at the forefront of intelligent public finance.

> *"Technology is only as effective as the governance framework supporting it. Responsible AI is what builds the public trust required for adoption at scale.*

In this sector

Department of Finance	p. 60
Dubai Government Human Resources Department	p. 61
GDRFAD	p. 62
Dubai Public Prosecution	p. 63
Judicial Council of Dubai	p. 64
Awqaf Dubai	p. 65
Mohammed Bin Rashid Housing Establishment	p. 66

REFERENCE [1] Department of Finance, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



1a DEPARTMENT OF FINANCE:

Intelligent Public Finance

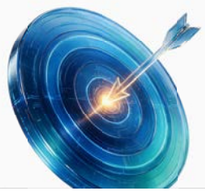
AILTIZAM, Dubai Finance's agentic AI platform, brings intelligent, bilingual governance, risk and compliance verification to public financial management across the department.



"At Dubai Finance, AI is not a tool we adopt but a pillar of how we safeguard Dubai's economic leadership and the long-term sustainability of its public finances."

H.E. Abdulrahman Saleh Al Saleh
Director General, Dubai Finance

AILTIZAM turns weeks of manual governance, risk and compliance checking into minutes, with every finding traceable to the law, and returns up to 1,860 hours a year
to Dubai Finance's professionals.



One scalable platform for governance, risk and compliance

AILTIZAM is an agentic AI platform for governance, risk and compliance, built to scale to all Dubai Government entities. It advances Dubai's AI Vision, the Dubai Universal Blueprint for AI and the D33 economic agenda by embedding AI in government operations, lifting productivity, and giving entities a trusted way to adopt it. Its core engine checks all 204 Dubai Finance processes, and a library of 300+ policies from multiple Dubai Finance departments, against 272 laws and 4,046 articles. Specialist AI agents identify gaps, recommend corrective actions, and compile bilingual Arabic and English reports in 5 to 15 minutes, down from 2 to 6 weeks for a process review and 5 to 25 days for a policy. Every finding is traced to its legal article, and all AI processing stays within the UAE.

Live use cases

Dubai Finance's Digital Payments contracts are now reviewed by AI agents on the new Contract Intelligence platform, cutting review from 10 to 15 days to 10 minutes. The Tax Operations Team's invoices follow the same route through Tax Pre-Filing Automation, screening 500,000+ invoices in just over 10 minutes and cutting turnaround from 7 weeks to 1.

A broader portfolio is in design and delivery: Public Private Partnerships, Public Procurement, Public Assets, Human Resources and Information Security Policy Compliance, all reusing the same secure, agentic core.

Outcomes and direction

Across the platform, 27 AI agents, a number set to grow with each new use case, carry the routine verification work, releasing 125 to 155 working hours of review capacity a month, some 1,500 to 1,860 hours a year, a figure that will rise as the future portfolio comes online. Dubai Finance professionals are freed for the judgement-intensive work that machines support but do not replace. AILTIZAM follows the Dubai AI Principles: sovereign hosting, human oversight and full auditability.

AILTIZAM turns governance, risk and compliance into a measurable, scalable capability: govern first, automate what governance has made trustworthy, and measure everything. Its next chapter is to extend that discipline to every department of Dubai Finance.

REFERENCE [1] Department of Finance, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



1b

DUBAI GOVERNMENT HUMAN RESOURCES DEPARTMENT:

The Workforce, Re-equipped

From voice-driven task management to AI-conducted candidate interviews, DGHR is rebuilding how government itself is staffed and run



"Our vision is centered on building an AI-enabled government HR ecosystem capable of anticipating workforce trends, strengthening institutional productivity, and empowering leadership with data-driven insights."

H.E. Abdulla Ali bin Zayed Al Falasi
Director General, Dubai Government Human Resources Department

"Talent Bot shortlists, analyses, and interviews candidates by voice at 85 per cent matching accuracy, returning thousands of recruiter hours to human judgement.

Live in production

DGHR's systems attack the two most time-hungry parts of government work: coordination and recruitment. Mahami, in production since June 2023, is a natural-language task management assistant that lets employees create and track work by voice, cutting task creation time by 80 per cent. Talent Bot, deployed in the fourth quarter of 2025, takes on the recruitment funnel end to end: it matches CVs to job descriptions, produces a strengths-and-weaknesses analysis of each candidate, conducts voice-based first interviews against AI-generated questions, and scores the shortlist, at 85 per cent matching accuracy. The effect is thousands of recruiter hours redirected from screening to selection, with a scaling path towards the Dubai Careers platform. A third system, the AI-powered military medical committee assistant deployed in the same quarter, streamlines medical committee workflows for military personnel, engineered to cut processing time by 70 per cent and remove five procedural steps.

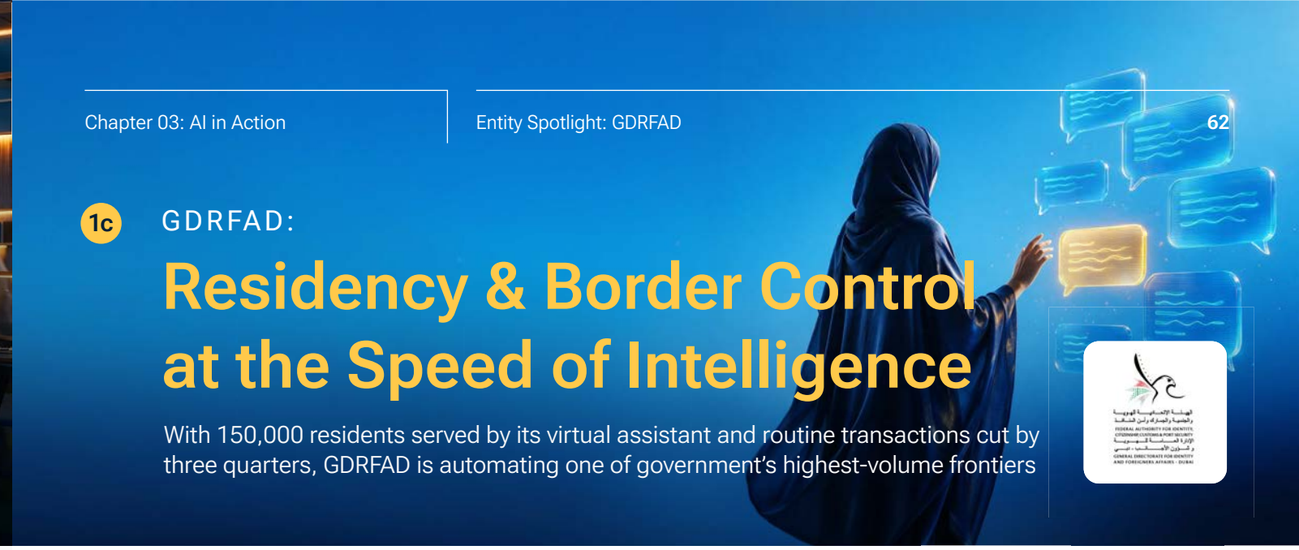
An ecosystem, not a toolset

The department builds with partners across the local technology sector: task intelligence with Cyberspace, talent matching with Diverge, an AI legislative portal with Brixbey that helps draft and interpret HR law and answer legal enquiries from precedent, and the reimagined Khbrati talent network with Neol. It is a working example of the open-innovation direction H.E. Al Falasi sets out: government functions enabled by AI-driven insight, data governance as the foundation of trust, and partnership as the accelerator.

Outcomes and direction

For DGHR the measure of success is human, not technical. Its stated conviction is that AI's value lies in expanding human potential rather than replacing it, which is why workforce readiness, structured change management, and responsible AI governance sit at the centre of the approach. As the entity that sets the standard for how the Government of Dubai hires, develops, and organizes its people, DGHR is positioning the government workforce itself as the emirate's decisive AI advantage. The next horizon is an HR ecosystem that anticipates workforce trends before they arrive, giving Dubai's leadership predictive sight of the skills the D33 agenda will demand.

REFERENCE
[1] Dubai Government Human Resources Department, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



1c

GDRFAD:

Residency & Border Control at the Speed of Intelligence

With 150,000 residents served by its virtual assistant and routine transactions cut by three quarters, GDRFAD is automating one of government's highest-volume frontiers



"GDRFAD places digital transformation at the core of its strategy, aiming to embed AI across all operational and service domains to deliver more efficient, flexible, and responsive public services."

H.E. Lieutenant General Mohammed Ahmed Al Marri
Director General, General Directorate of Identity and Foreigners Affairs, Dubai

AED 5.81B
projected cost reduction by 2035 from replacing traditional counters and Smart Gates with the Red Carpet service
Red Carpet, live since 2025

Live in production

GDRFAD holds ISO 42001 certification for its AI management system. Salama, its AI virtual assistant launched in 2024 and now integrated into Dubai AI, is the front door to residency services: 150,000 active users, 40 per cent of digital service users, growth above 20 per cent, and transactions completed without a centre visit, from residency applications to violation payments and entry permit extensions, at 88 per cent model accuracy and 99.9 per cent uptime. Behind it, the Proactive Intelligence System, live since 2024, automates visa and residency transactions across entry permits, ports, establishments and labour relations; serving an estimated 250,000 users a year, it cuts turnaround on routine transactions by 75 per cent. A generative AI sentiment system, in production since 2023, reads public feedback across social channels at an estimated 85 per cent accuracy, and the GDRFAD Data Lake, launched in 2026, consolidates service data as the foundation for predictive services.

At the border, the Red Carpet service, introduced in 2025, clears ten passengers at a time without stopping or presenting documents, in an average of 3.4 seconds. Since launch it has served 1.3 million passengers at an 89 per cent eligibility rate, with usage up 312 per cent, seven times the capacity of traditional Smart Gates, 99.8 per cent passenger happiness and 100 per cent partner happiness.

"Salama, GDRFAD's AI assistant embedded in Dubai AI, serves 150,000 users around the clock at 88 per cent model accuracy, answering questions and completing transactions.

Outcomes and direction

For citizens and residents, the outcome is a border control and residency system that moves at the speed of the city it serves: faster approvals, fewer visits, and services that anticipate rather than react. GDRFAD's AI Strategy for 2024 to 2030 frames this as a permanent operating model, pairing operational excellence and data-driven decision-making with strict standards for privacy, security, and ethical use, and backing it with training programs, innovation challenges, and partnerships with leading global technology companies. As the authority through which millions encounter Dubai first, GDRFAD is defining what AI-era government feels like at first contact. Its next step is to extend the proactive model until the routine transaction disappears altogether, leaving human officers for the cases that truly need them.

REFERENCE
[1] General Directorate of Identity and Foreigners Affairs, Dubai, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



1d

DUBAI PUBLIC PROSECUTION:

Justice, Digitally Grounded

A dedicated AI team under the Attorney General is building intelligence into criminal case management on foundations of governance and trust



"Recognizing the strategic importance of AI technologies, the organization is committed to integrating these solutions within the judicial ecosystem to enhance operational efficiency in criminal case management, support data-driven decision-making, and improve the quality of services provided by Public Prosecution, while upholding the principles of justice, transparency, and legal integrity."

H.E. Counsellor Essam Issa Al Humaidan
Attorney General of Dubai

> *"In judicial work, the sequence is the safeguard: governed data, secure platforms, and human accountability come first, and intelligence is built on top of them.*

Foundations in production

Between 2023 and 2025 Dubai Public Prosecution strengthened its digital core to carry intelligence: the digital criminal case management systems through which prosecutions are run, data analytics, and electronic services for prosecutors, employees, and all service users. These production systems generate the governed, structured case data on which trustworthy AI depends. Around that core sit robust data governance, secure digital platforms, and partnerships across government. A dedicated Artificial Intelligence Team, led directly by the Attorney General, develops and implements AI initiatives in line with Dubai's digital transformation strategy.

Two AI use cases

The first is an AI assistant for the penal order, the decision by which the Public Prosecution settles a misdemeanour or contravention with a fine where the law does not require imprisonment. The assistant analyses the case file, checks that the conditions are met, including the type of charge and whether the accused has confessed, and drafts the decision with the fine the Attorney General has set for the charge. The prosecutor reviews a summary of the case and its evidence and approves the decision. The accused can accept by text message and pay through a link, or object; payment lifts restrictions such as a travel ban at once. Available around the clock, including public holidays, the service is designed to settle more cases faster for every service user, tourists among them.

The second, Compliance and Competition, applies AI to prosecutors' performance data to raise productivity and compliance with procedures. Each prosecutor receives a live percentage score on three criteria: investigation (such as investigating before recording a disposition), disposition (meeting time limits and acting promptly on penal orders), and compliance with regulations (attendance, training plans, and response times to service users). Prosecutors and senior leadership see each score beside those of colleagues, to motivate through competition.

Outcomes and direction

Both use cases serve the Public Prosecution's strategic objective of just and efficient investigation and prosecution, and both are measured by the speed of investigation and disposition. Because the prosecutor approves every penal order the assistant drafts, intelligence extends the reach of justice without loosening its standards.

REFERENCE [1] Dubai Public Prosecution, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



1e

JUDICIAL COUNCIL OF DUBAI:

Building the Bench for AI

The General Secretariat of the Judicial Council of Dubai is digitizing its services, so that intelligence arrives on foundations built to hold it.



"We have seen in recent years a major shift in technology and a rush to benefit from AI-driven projects, which would be the foundation for future tech projects and provide a more adaptive approach to utilizing technology in providing user friendly services to all."

H.E. Prof. Abdulla Saif Alsuboosi
General Secretary of the Judicial Council of Dubai



A newly established entity can learn from every AI initiative that came before it. The Judicial Council is doing exactly that, by design.

Digital foundations in production

The General Secretariat of the Judicial Council of Dubai, restructured in 2021 to provide administrative, technical, and legal support to Dubai's Judicial Council, has spent its first years building the digital estate on which judicial coordination now runs. The Council's public website, launched in 2024 at tjc.ae, provides the public with information on Dubai's judicial bodies, including the Annual Report of achievements across the judicial entities. Alongside it, an internal platform delivers multiple online services to judicial members, continuously reviewed and refined. A meetings platform now extends this estate, digitizing the scheduling of Judicial Council sessions and giving members a secure space to discuss matters pending decision.

Capability across the judicial system

The Council's ecosystem is investing in people as well as in platforms. The General Secretariat of the Judicial Council of Dubai has made comprehensive digital transformation and the employment of AI a core objective of its future strategy: its entire workforce has completed AI qualification programs, applies AI tools in day-to-day tasks, and has completed the One Million AI Prompters program, with an advanced course to follow this year. Its operational plan directs that capability towards case review support, gap analysis, and, in time, an intelligent virtual partner, under a vision of a smart, proactive, and sustainable judicial entity.

Outcomes and direction

The General Secretariat of the Judicial Council of Dubai is candid about its sequencing, and that is its strength. Its stated vision is to introduce AI into its platforms once the digital foundations are fully bedded in: first to summarize the matters brought before Council members, speeding decisions and raising the quality of submissions, then to support preliminary study of cases and evaluation of legislation. Discussions with its cloud and service partners at MORO are already shaping how that intelligence will be delivered, and as a young entity it deliberately draws on the lessons of Dubai's more established AI programs. In a domain where trust is the entire product, the Judicial Council of Dubai is showing what responsible adoption looks like from day one: govern, digitize, test, train, and only then automate.

REFERENCE [1] Judicial Council of Dubai, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



1f

AWQAF DUBAI:

Intelligence in the Service of Endowment

Dubai's endowments authority has rebuilt its data foundations and trained its entire workforce, bringing AI into the stewardship of community assets

> *"At Awqaf Dubai, we see artificial intelligence not as a mere operational tool, but as vital infrastructure that contributes to improving the management of endowment assets, enables precise decisions, and maximises returns in a way that serves the community effectively and continuously."*



H.E. Ali Mohammad AlMuttawa
Secretary General of Awqaf Dubai

From exploration to operations

The past three years brought rapid, qualitative advances in AI, and Awqaf Dubai has translated them into new capability for one of the emirate's oldest institutions: the endowment.
Over this period the organization's efforts matured from early exploration into practical applications inside its vital operations, supported by a flexible institutional culture and a commitment to sustainable innovation. These technologies now support decision-making across the endowment portfolio, improving the management of waqf assets and strengthening the sustainability of the resources that fund community benefit.

Data first, people always

The organization's central insight is disarmingly simple: no advanced AI system can deliver without accurate, integrated, continuously updated data. Acting on it, Awqaf Dubai restructured its endowment data repositories into a unified, analysable foundation, enabling it to improve investment performance, forecast returns, and direct waqf resources to community initiatives with efficiency and transparency.
The human investment ran in parallel: every employee has received comprehensive training that goes beyond using AI tools to critically evaluating their outputs and assuring their reliability, reinforced by a culture of digital security and institutional governance that protects sensitive endowment data.

Explainability as obligation

Endowment work carries a distinctive duty: every investment decision and every allocation of returns must be justifiable to the community on whose behalf assets are held. Awqaf Dubai treated the explainability of AI as a principal challenge and answered it directly, developing mechanisms that allow decisions to be followed transparently, strengthening trust and guaranteeing accountability before the public. It is responsible adoption shaped by the mandate itself.

> *"Every endowment decision must be explainable to the community it serves. Awqaf Dubai built that requirement into its AI from the start.*

Direction

The path ahead integrates advanced AI systems seamlessly into endowment operations, delivering precise analytics and predictive reporting that sharpen strategic decisions, improve asset investment, and deepen the financial sustainability of the waqf. The stated ambition is that every endowment will carry a strategic plan supported by AI, maximizing its impact and securing the continuity of giving for the community. As steward of the endowment sector, Awqaf Dubai is demonstrating that the AI era and the waqf tradition strengthen one another: intelligence extends the reach of giving, and giving gives intelligence its purpose.

REFERENCE [1] Awqaf Dubai, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



1g

MOHAMMED BIN RASHID HOUSING ESTABLISHMENT:

AI at the Heart of Housing Transformation

AI-powered decisions process housing applications instantly, while generative AI transforms family aspirations into home designs in seconds



> *"At the heart of this transformation is the strategic importance of data. Artificial intelligence delivers meaningful impact only when supported by high-quality, well-governed, and integrated data ecosystems."*

H.E. Mohammed Hasan Alshehhi
Chief Executive Officer, Mohammed Bin Rashid Housing Establishment

SMART HOUSING DESIGNER

BEFORE 90 days → NOW INSTANT design and BoQ

cut in housing design time

98.75%
customer happiness measured

Live in production

The Housing Intelligence Matrix, in production since 2021, is the establishment's AI-powered decision-making platform, enabling instant processing of eligible housing applications across 12 services. Operating at 95 per cent accuracy with 100 per cent service adoption, the platform has reduced manual employee effort by 30 per cent, saving approximately 9,000 working hours and generating AED 2.3 million in cost savings, while contributing to an unprecedented customer happiness rating of 98.75 per cent.

In the fourth quarter of 2025, MBRHE added two customer-facing solutions. The MBRHE AI Assistant is a bilingual Arabic and English conversational service available on WhatsApp, the web, iOS, and Android. Powered by retrieval-augmented generation and grounded in the establishment's policy documents, it provides around-the-clock guidance on eligibility, documentation, and next steps, with full human auditability and seamless hand-off to officers for final decisions. It has saved more than 1,000 service hours and AED 2 million in costs.

The Smart Housing Designer uses generative AI to turn a family's preferences into a complete home design in seconds: 2D and 3D plans, virtual walkthroughs, facades, editable AutoCAD files, and bills of quantities with estimated costs. In four months it drew 1,100 customers and generated estimated savings of AED 255,000.

Outcomes and direction

For Emirati families, the impact is tangible: faster housing decisions, round-the-clock bilingual assistance, and personalized home designs brought to life in seconds. Beyond customer services, MBRHE is introducing a unified, bilingual AI workspace of intelligent assistants and agents to enhance employee productivity, collaboration, and knowledge sharing.

Strong governance supports the transformation: a dedicated AI Committee and an Artificial Intelligence Policy aligned with the Dubai AI Policy and global best practice, with data protection, controlled access, transparency, traceability, human oversight, and accountability maintained throughout. MBRHE next aims to advance predictive and agentic AI to anticipate citizens' housing needs, enable proactive services, and support strategic housing planning across the emirate.

REFERENCE [1] Mohammed Bin Rashid Housing Establishment, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.

2

MOBILITY & URBAN DEVELOPMENT:

The City That Learns to Move

From the street grid to the skies above it, Dubai's mobility and urban institutions are running AI in daily production, and planning the city around what it learns



A reflection by H.E. Mattar Al Tayer, Director General, Chairman of the Board of Executive Directors, Roads and Transport Authority

Mobility is where a city's ambition becomes a daily, lived experience. At the Roads and Transport Authority, our updated Artificial Intelligence Strategy is built on a clear vision: to reinforce RTA's global leadership in AI-powered mobility, creating an integrated, self-optimizing transport ecosystem that enhances quality of life for everyone in Dubai. A foundational stage, from 2017 to 2021, gave us smart decision support, an enterprise command and control centre, intelligent traffic systems and the Smart Yard for driver testing; a localized stage, from 2022 to 2025, put an upgraded Mahboub, an enterprise AI platform, smart road-monitoring vehicles, smart monitoring centres and driver-monitoring systems into daily production; and in 2026 we entered the transformational stage, an updated AI Strategy for 2026–2028, implemented over two and a half years, which rewires 41 business domains around six purpose-built AI platforms and 200 AI products and targets more than AED 300 million in annual added value.

From data to decision

None of this is possible without disciplined foundations. Since 2017, RTA has built a Big Data Platform governed by a strict data framework and a culture that treats every byte as a strategic asset; it is now being operated alongside the Enterprise AI Platform. In May 2025, RTA established the AI and Data Science Centre as the single owner of AI strategy and delivery across RTA. In February 2026 an AI Strategy Committee and domain-level execution squads were formed around it, giving every agency a sprint-based rhythm for identifying, building and operating AI products, with AI champions inside each domain keeping innovation rooted in operational reality.

The past three years have taught us that technology is secondary to data quality and human readiness. Our 2026–2028 strategy therefore replaces separate products with six reusable platforms shared by every domain, wave by wave. Seventy per cent of our employees are already enabled on Copilot, and we expect half of RTA's operations and services to run through an agentic AI ecosystem within two and a half years.

> *"Technology is secondary to data quality and human readiness. Sustainable production requires a mindset shift from project to product.*

In this sector

Roads and Transport Authority p. 68
Runs the intelligence of the street: traffic, licensing, public transport and the road assets beneath them.

Dubai Municipality p. 69
Orchestrates the urban fabric itself, from a city-scale Digital Twin that mirrors every building and asset to AI-assisted cleanliness monitoring on the streets.

Dubai Civil Aviation Authority p. 70
Governs the airspace above, where AI now approves aviation services in real time and supervises drone operations across the emirate.

Three mandates, one direction: a city that senses, learns and responds as a single system.

REFERENCE
[1] Roads and Transport Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.
[2] Dubai Municipality, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.
[3] Dubai Civil Aviation Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.

2a

ROADS AND TRANSPORT AUTHORITY:

Intelligence on Every Journey

From a dedicated AI and Data Science Centre to six purpose-built agentic platforms, the Roads and Transport Authority is running AI across the mobility system, in production, at scale, and is now rewiring every domain around it



"The strategy marks an advanced stage in RTA's digital transformation journey, paving the way for a fundamental shift from AI applications that support operations to intelligent systems capable of analysing data, predicting outcomes, supporting decision-making and delivering end-to-end services."

H.E. Mattar Al Tayer
Director General, Chairman of the Board of Executive Directors, Roads and Transport Authority

AI STRATEGY 2026–2028
SIX PLATFORMS

BEFORE (2022–2025)		NOW (2026–2028)
1	→	6
product at a time, built domain by domain		platforms, shared across 41 domains

50%	AED 300M+
annual added value targeted	of operations and services on agentic AI within 2½ years

The AI and Data Science Centre's flagship deliverable is the updated AI Strategy 2026–2028, launched on 1 November 2026: RTA's move from AI products built one domain at a time to a transformational stage. Over two and a half years it equips 41 business domains with 200 AI products and more than 1,000 enterprise agents, and targets more than AED 300 million in annual added value, up to 5 per cent less traffic congestion, up to 30 per cent faster case and complaint handling, up to 20 per cent faster service at Customer Happiness Centres and up to 20 per cent fewer accidents among newly licensed drivers.

Six purpose-built platforms carry the strategy. Planning and multimodal operations (six domains): bus, marine and rail planning and monitoring; a line review that took three days takes two minutes, and the agentic Bus Network Planner assesses five times faster. Mahboub, agentic services and customer care (six domains): services completed inside the conversation on RTA's app, website and WhatsApp; five vehicle licensing services already close end to end in under five minutes through UAE PASS and Dubai Pay. Congestion management and prediction (four domains): iTraffic, UTC-UX Fusion signals and smart pedestrian signals, targeting 5 million driver hours saved a year. Asset inspection and management (three domains): LiDAR surveys 80 kilometres of road assets a day at 97 per cent accuracy, and an agentic No Objection Certificate system is set to review some 130,000 right-of-way applications a year three times faster, saving more than 700,000 days annually. Smart monitoring and digital twins (six domains): the Smart Monitoring Centre already performs 46 per cent of passenger-transport inspections. Corporate services (16 domains): Copilot, at active usage above 90 per cent and a 10 per cent saving in working time.

Delivery runs in six-month waves. The first two, from April 2026, built a portfolio of 40 AI products and launched 11, completed four foundational projects in cloud, big data, AI infrastructure and data governance, and enabled half of RTA's employees on Copilot and Agent Builder; 70 per cent are now active on it, returning 10 to 15 times the licence cost.

Together these systems position RTA as the engine of AI-powered mobility in Dubai. With 25 per cent of trips targeted for autonomous mode by 2030, robotaxis, flying taxis and an agentic journey planner on the roadmap, and a Zero-Service RTA in 2030 as the destination, RTA is building the mobility system the intelligent city will run on.

REFERENCE
[1] Roads and Transport Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



2b DUBAI MUNICIPALITY:

The City's Cognitive Nervous System

A city-scale Digital Twin, AI-powered field enforcement and intelligent public health screening: Dubai Municipality is orchestrating the urban environment with production AI



"We view AI not as an isolated technological tool, but as the foundational nervous system of a world-class city. Our goal is to leverage intelligent systems to create a 'Human-Centric Smart City' ... to enhance the quality of life for every resident and visitor."

H.E. Eng. Marwan Ahmed Bin Ghalita
Director General, Dubai Municipality

AI is a cultural shift as much as a technical one, and intelligence is only as strong as the data ecosystem beneath it.



Dubai Municipality is evolving from a traditional service provider into a cognitive urban orchestrator: an institution whose operations are predictive, whose infrastructure is resilient, and whose services adapt to the people who use them. Its unified intelligence architecture concentrates AI where the city feels it most, in sustainable urban planning, intelligent field enforcement and digital twin infrastructure.

The flagship is the Dubai Digital Twin, launched in 2024 and expanded in 2026 into an integrated platform spanning more than 195,000 buildings, 330,000 public facilities and 280,000 infrastructure assets: a living digital replica of the emirate that spans master plans, roads, trees, green spaces and utility lines down to the details of individual residential units, refreshed with real-time data. Government entities and urban planners use it to review sites virtually rather than physically, simulate rainfall and future scenarios, and manage infrastructure and assets from a single environment. Next in line is an AI agent, now in development, that will check villa designs against the Dubai Building Code and issue permits automatically.

In the field, the Eltizam application, introduced in 2025, gives municipal inspectors smart scanning that reads Emirates IDs and vehicle plates instantly, replacing manual data entry with near-instant issuance of violation records across eight categories of public cleanliness offences. AI-powered cameras mounted on waste collection vehicles extend that reach, using computer vision and machine learning during routine rounds, while the RASID platform tracks the waste fleet by satellite around the clock, analysing GPS data in real time to keep operations on course.

Public health screening is intelligent too. The Montaji Plus system uses AI to register and classify food and consumer products and verify their compliance, sharply reducing time-to-market for the more than 1.5 million products on its books while strengthening compliance. In the city's parks, the Smart Track system applies computer vision to running tracks, helping runners monitor their performance while giving park managers live insight into usage, safety and peak demand.

The common thread is a hard-won lesson: AI is a cultural shift as much as a technical one, and intelligence is only as strong as the data ecosystem beneath it. Dubai Municipality has answered with rigorous data governance, explainable AI to build trust, and sustained internal upskilling. As the Municipality moves towards autonomous municipal services, it is building a city that thinks and responds in real time, and setting the global benchmark for how a metropolis is run.

REFERENCE [1] Dubai Municipality, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



2c DUBAI CIVIL AVIATION AUTHORITY:

Anticipatory Governance for Intelligent Skies

Real-time AI approvals, drone airspace supervision and a governance-first operating model make the Dubai Civil Aviation Authority a template for AI in safety-critical regulation

"Our vision extends beyond automation. We are building a proactive, data-driven regulatory environment that enhances aviation safety, strengthens compliance assurance, and supports evidence-based policy development."



H.E. Mohammed Abdulla Lengawi
Director General, Dubai Civil Aviation Authority

In aviation, trust is the product. The Dubai Civil Aviation Authority has established an AI strategy aligned with its overall entity strategy 2025-2028 and delivered under a governance-first model, in step with Dubai's digital transformation agenda and the long-term ambitions of UAE Centennial 2071. A formal AI workforce, established in 2025, and comprehensive lifecycle frameworks embed transparency, accountability, fairness and cybersecurity resilience at every stage of deployment. Innovation scales because it is governed, a principle that matters most in a safety-critical ecosystem, and one that has turned AI from a promising initiative into an institutionalized capability.

The results are measurable. The AI Instant Approval Engine, live since 2021 on the DCAA digital portal, validates aviation service requests against regulatory rules, operational parameters and safety conditions, has processed more than 83,000 automated transactions. Where manual approvals once took an average of 1 working day across thousands of cases, decisions now arrive in real time, a zero-day turnaround. A structured 2 per cent verification sample of 113 records confirmed 100 per cent decision accuracy, and customer satisfaction exceeds 95 per cent, with users describing the platform as a "super system".

Above the city, the Dubai Monitoring System supervises drone operations across Dubai airspace in real time. The AI-enabled platform analyses telemetry, flight paths and weather conditions to detect risks such as restricted-zone violations, alerting operators and authorities automatically. It has autonomously processed 24,214 successful flights and anchors the emirate's readiness for advanced air mobility. On the customer side, an AI-powered WhatsApp service launched in November 2025 answers regulatory enquiries around the clock, drawing on verified knowledge bases and escalating complex cases to human agents; satisfaction stands above 96 per cent.

External validation has followed: the Global Artificial Intelligence Award in the Silver Category and the UAE Innovation Award in the Gold Category, both received in 2025; the Best AAM Infrastructure Project at the AAM Awards 2025 Asia in Hong Kong; and the Transport Creativity Award in the Innovation Category at the Dubai Sustainable Transport Award in 2026. Together they recognize not just the technology but the governance and measurable public value behind it. DCAA describes its shift as moving from reactive control to anticipatory governance, and the description fits: the regulator now detects risk before it materializes and approves compliant activity the moment it is requested.

Looking towards UAE Centennial 2071, DCAA is advancing predictive safety intelligence, digital twin integration and scalable ecosystem monitoring, positioning Dubai's skies, and the authority that governs them, at the forefront of intelligent aviation regulation worldwide.

REFERENCE [1] Dubai Civil Aviation Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.
[2] Dubai Civil Aviation Authority, "Director General's message on artificial intelligence," State of AI Report 2026 entity submission (supplementary), Dubai, UAE, 2026.

3

ENERGY & UTILITIES:

The AI-Native Utility

Dubai's power and water system is being re-engineered end to end by artificial intelligence, from the turbine hall to the customer conversation



A reflection by H.E. Saeed Mohammed Al Tayer
MD & CEO, Dubai Electricity and Water Authority

The forward-looking vision and directives of the UAE's wise leadership underpin Dubai's transformation into a global hub for sustainability and artificial intelligence. Under the vision of His Highness Sheikh Mohammed bin Rashid Al Maktoum, Vice President and Prime Minister of the UAE and Ruler of Dubai, and the directives of His Highness Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, Crown Prince of Dubai, Deputy Prime Minister and Minister of Defence of the UAE, and Chairman of The Executive Council of Dubai, AI has been placed at the heart of future-oriented initiatives designed to serve humanity.

At the strategic level, for Dubai Electricity and Water Authority (DEWA), this is not a matter of launching smart applications. Our AI journey is a comprehensive strategic transformation that redefines the role of utilities, encompassing the re-engineering of the entire value chain: the planning, generation, transmission and distribution of electricity and water, through to the delivery of advanced digital services.

Proof at every layer

At the operational level, DEWA developed the world's first AI-based intelligent controller for gas turbines used in power generation and water desalination, enabling autonomous turbine control, and has extended this capability across the combined cycle power block. At the Mohammed bin Rashid Al Maktoum Solar Park, smart robots, automated cleaning, solar tracking and production forecasting enhance efficiency and reliability. In services, Rammas, our AI-powered virtual employee, has handled more than 13.2 million customer interactions since 2017 without human intervention, and DEWA became the first organization in the UAE to receive ISO certification for trustworthiness in AI. We were also the first government entity worldwide to provide services through the ChatGPT platform, and we are in the process of expanding the offering to 21 services.

The establishment of Digital DEWA (MORO, InfraX and DigitalX), our digital arm, marks the shift from selective technology adoption to a fully integrated digital ecosystem, guided by an ambitious roadmap to position DEWA as the world's first AI-native utility, in line with the UAE National Artificial Intelligence Strategy 2031.

> *"AI is no longer a technological option. It is a strategic imperative that is reshaping economies, redefining services and transforming the daily lives of societies.*

AI strengthens economic resilience by reducing costs, boosting productivity, attracting investment and creating high-value opportunities for national talent. DEWA's AI ecosystem brings together Microsoft, Google, OpenAI, SAP, GitHub and Siemens Energy, alongside specialized partners in energy analytics, digital twins, solar intelligence and agentic platforms. Through Digital DEWA, this ecosystem forms an integrated operating model that accelerates innovation and advances the UAE's knowledge based economy.

In this sector

Dubai Electricity and Water Authority pp. 72 to 75
Runs production AI across four layers: the customer conversation, the power plant, the grid, and the partner ecosystem around them.

In Dubai, one institution carries this sector end to end.

DEWA's Agentic AI Transformation

100%
DEWA's Services are now Agentic

Dubai Electricity and Water Authority (DEWA) achieved 100% transformation of its service portfolio, converting all 21 services into agentic AI-powered services in just one-quarter of the government-mandated timeframe for 50% transformation. This milestone reflects a fundamental shift in service delivery through streamlined procedures, simplified requirements, task automation, faster processing, and enhanced system integration. The transformation covers electricity and water supply services, technical notifications, and refund, consumption and payment reconciliation, extending agentic AI capabilities across DEWA's services and customer journeys.

75%
Faster Move-in

Reducing the time required to apply for electricity and water activation of supply by 75%

For the electricity and water activation of supply service, which processes approximately 210,000 transactions annually, the transformation reduced the required customer inputs from 10 to just 3, cut customer journey steps from 9 to 3 and reduced transaction completion time.

24/7
Payment Anomaly Monitoring

Payment reconciliation cycles cut by 95%

Payment reconciliation supports 13 DEWA services and processes about 13.3 million transactions annually. The implementation of agentic AI has transformed the reconciliation process into an integrated ecosystem capable of processing and validating data, detecting anomalies, managing exceptions and continuously monitoring operations 24/7. As a result, the daily reconciliation cycle has been reduced from around 4 hours to just 12 minutes, representing a 95 percent improvement in processing time.

94.5%
Faster in document stamping

DEWA Accelerates Building and Demolition NOC Services through AI-Powered Document Verification and Digital Stamping

The AI Powered NOC Services enhance Infrastructure, Demolition and Building NOC services by automating document validation and digital stamping, reducing manual effort while retaining final decisions with authorised DEWA users. The solutions accelerate processing, improve consistency and enable engineers to focus on higher-value technical reviews. Document Verification is expected to reduce review time by 60% per application, while Digital Stamping reduces processing from 6 steps to 1 and from 4.52 minutes to approximately 0.25 minutes. These improvements enhance efficiency.

DEWA Harnesses AI to Accelerate Refund Services and Streamline Consumption Verification

The Refund Service AI Agent streamlines the refund journey, reduces manual effort and improves the customer experience. Consumption Investigation is expected to reduce technician preparationtime for field visits by 50%, from 6 to 3 minutes, and increase throughput by 20%. The Conversation Agent provides instant responses to customer notifications, eliminating associated waiting time and manual customer calls. Together, these capabilities support faster service and enable backoffice teams to focus on complex cases, while maintaining compliance and service continuity.

Response to technical notifications Services

For the smart handling of electricity and water technical notifications, DEWA redesigned the case assignment process through an intelligent ecosystem that receives, classifies and verifies requests, and coordinates workflows through to resolution. The assignment process has evolved from relying on 16 coordinators to fully autonomous execution around the clock.

These results reaffirm that DEWA's adoption of agentic AI is guided by clear and measurable objectives focused on accelerating service delivery, streamlining procedures, enhancing operational efficiency and improving service continuity and quality. This supports Dubai's vision of consolidating its position as a global leader in the deployment of AI and future technologies.

3a DUBAI ELECTRICITY AND WATER AUTHORITY:

Sixty-Eight Ways to Run a Utility on Intelligence

With production AI in the turbine hall, on the grid and inside 13.2 million customer conversations, DEWA is the working model of the AI-native utility



"DEWA's AI journey represents a comprehensive strategic transformation that redefines the role of utilities ... it encompasses the re-engineering of the entire value chain, from the planning, generation, transmission, and distribution of electricity and water to the delivery of advanced digital services."

H.E. Saeed Mohammed Al Tayer
MD & CEO,
Dubai Electricity and Water Authority

28,530 tn
of CO2 emissions avoided

AED 9.34M
Cost Saving

Dubai Electricity and Water Authority's contribution to this report showcases a broad AI portfolio, with 68 applications in production spanning every layer of the utility.

Rammas, the Smart virtual employee

When DEWA launched Rammas in 2017, it became the first government entity in the UAE to deploy AI for customer interaction at scale. Today the generative AI-powered virtual employee converses around the clock in Arabic and English, and the numbers describe a workforce in itself: more than 13.2 million user questions handled across over 4 million sessions, satisfaction of around 90 percent based on the annual Customer Satisfaction Survey, model accuracy of 97.54 percent and availability of 99.97 percent. Rammas saves hundreds of thousands of working man-hours every year, 344,231 hours in 2025 alone, translating into annual savings of AED 12 to 14 million in each of the past three years, while DEWA's paperless operation has saved the equivalent of 57,000 trees and cut 15,000 tonnes of CO2.

In January 2026, DEWA extended this lead by becoming the first utility, and the first government entity globally, to launch official services on the OpenAI Apps Directory, moving customers from clicking and tapping to simply asking. The services delivered through the new channel have operated with 100 percent accuracy and uptime, an early marker of where government service delivery is heading.

Intelligence inside the power plant

At the heart of gas turbine operations, AI is enabling autonomous, real-time performance optimization. The Gas Turbine Intelligent Controller (GTIC), a world-first solution developed by DEWA with Siemens Energy, uses AI and Machine Learning to continuously learn from operating conditions and dynamically optimize turbine parameters. Directly interfacing with the turbine control system, GTIC recovers lost generation capacity, improves thermal efficiency, reduces fuel consumption and lowers emissions. The solution delivers AED 9.33 million in cost savings, with a 4.6-year ROI and no run cost. GTIC also reduces CO_2 emissions by 28,530 tonnes and NOx emissions by 22.9 tonnes annually, positioning DEWA as a global pioneer in intelligent, efficient and sustainable power generation.

An ecosystem, not a project

What distinguishes DEWA is the breadth of the ecosystem behind these deployments. Its private-sector AI bench includes Microsoft, Google, OpenAI, SAP, GitHub and Siemens Energy alongside specialists such as Bidgely for consumption analytics, Fugro for AI-powered infrastructure digital twins, Solargis and Solcast for solar intelligence and AIQ for agentic energy platforms. Through Digital DEWA, its digital arm, the authority has converted this network into an operating model rather than a supplier list. And the portfolio keeps widening. Earth observation AI now watches Dubai's coastline for oil spills and red tides, forecasts fog three days ahead to protect operations, and monitors construction progress from satellite imagery. On the water network, anomaly detection and proactive highwater-usage alerts help customers find leaks and conserve before the bill arrives, while a Smart Outage Management System sharpens the response when the grid needs it.

DEWA's ambition to become the world's first AI-native utility is increasingly reflected in its achievements: production AI across the value chain, world-first innovations in operations and customer services, and ISO certification for AI trustworthiness. At the Mohammed bin Rashid Al Maktoum Solar Park, strategic collaborations with international developers and technology providers further advance AI-driven operations, automated cleaning, forecasting, battery systems and power-conversion technologies.

AI on the grid

Between the plant and the customer sits a grid that increasingly thinks for itself. DEWASOL, DEWA's proprietary satellite-based model for solar irradiance estimation and forecasting, delivers 15-minute resolution irradiance maps across the UAE and feeds the Smart Distribution Control Centre's dispatch decisions for more than 100,000 Shams Dubai rooftop solar installations. Its companion, DEWACAST, applies deep learning to photovoltaic power forecasting, cutting forecast error for direct sunlight from roughly 76 percent to 23 percent and delivering more than 50 percent overall accuracy improvement, so the grid can absorb more renewable energy with confidence. And in the Distribution Control Centre, a generative AI Switching Generator, live since January 2025 and used daily by 60 controllers, drafts electrical switching plans that previously took more than an hour of manual coordination, operating at 97 percent accuracy and saving over 896 working hours a year, with 100 percent adoption among its users. Together, the forecasting systems alone contribute AED 7.34 million in annual savings through optimized solar grid integration and reduced curtailment, and the whole layer supports DEWA's target of 100 percent cleanenergy by 2050.

Plant Intelligent Controller (PIC): AI that autonomously optimizes the power block.



Plant Intelligent Controller (PIC)

At the heart of the power block, AI is enabling autonomous, plant-wide optimization. The Plant Intelligent Controller (PIC) is a digital twin-based, closed-loop controller that orchestrates two Gas Turbines, two HRSGs and one Steam Turbine as one integrated system, reducing reliance on manual decisions and asset-level optimization. With no human interaction required, PIC sends outputs directly to the control system to optimize efficiency under changing operating conditions. Online power-block testing confirms reliable outputs with a 0% error rate. PIC delivers AED 32.42 million in cost savings, with a 1.19-year ROI and nil run cost. It also reduces CO_2 emissions by 105,000 tonnes annually and is 100% compliant with the Digital Dubai AI assessment tool, advancing efficient, reliable and sustainable generation.

Every compliant invoice checked, posted and handed to payment without manual chasing



Supplier Invoice Verification Using AI

DEWA's AI-powered LPO Payments Agent is transforming supplier payment processing through intelligent automation, enhancing efficiency, accuracy and transparency. The agent automatically processes supplier invoices, extracts and validates key information, identifies discrepancies and directs exceptions to the relevant teams for resolution. Validated invoices are seamlessly posted in SAP and prepared for integration with the host-to-host (H2H) payment system.

The solution currently automates 85% of invoice transactions, reducing average processing time by 67%, from 3 minutes to just 1 minute per invoice. With human oversight maintained for exceptions, the platform accelerates payment processing, strengthens financial governance and ensures full traceability, supporting a more efficient, reliable and transparent supplier payment experience.



Smart Work Clearance System

Smart Transformation of the Work Clearance Application System

The Work Clearance Application process for periodic station visits has been transformed from a manual, effort-intensive workflow into an integrated digital journey for staff. The solution automates application creation, data entry, document processing, status tracking, reporting, and notifications. By reducing repetitive tasks and manual intervention, it helps accelerate processing, improve data accuracy and reliability, and strengthen operational visibility. The initiative enhances staff efficiency and service quality while advancing sustainable digital transformation.

DEWA's AI Adoption Journey

99.6%
Copilot Adoption

DEWA's AI Adoption journey began in 2024 with the ambition of becoming the first AI-Native organization, where every employee, role, and process is AI-enabled. Recognizing that access to AI tools alone would not ensure lasting adoption or organizational capability, DEWA expanded Microsoft 365 Copilot from 500 pilot licences in 2024 to 9,250 licences in 2025, supported by structured change management, AI governance, and a network of AI Champions, Ambassadors, and internal trainers.

96.6%
Satisfaction rate

What began as an AI awareness initiative evolved into an enterprise-wide capability-building programme, delivering more than 450 training sessions, reaching over 7,500 learners, accumulating 18,041 learning hours, and achieving 96.60% learner satisfaction by 2026.

The Adoption further broadened access through self-paced learning pathways featuring 223 lectures across two structured phases, engaging 1,320 active learners, complemented by hands-on workshops, webinars, knowledge sharing sessions, and practical experimentation.

7200
Trained in Agentic AI

AI Coder and AI Engineer learning tracks, hackathons, and agent-building activities helped employees progress from learning about AI to creating solutions for real business challenges, contributing to 99.60% Microsoft 365 Copilot adoption, 6,961 employees trained in Agentic AI, over 1,140 AI creators, and 6,188 AI agents created.

Building on these achievements, DEWA has designed 18 specialized AI Schools to establish structured, role-based pathways that progressively upskill, reskill, and certify employees, transforming the Adoption efforts into a sustainable workforce transformation engine

Through this next chapter, employees will collaborate with AI agents, redesign business processes, and create measurable organizational value, advancing the vision of an AI-Native DEWA where human and artificial intelligence work together to drive operational excellence, continuous innovation, and sustainable transformation across every role and division.

REFERENCE
[1] Dubai Electricity and Water Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.

4

HEALTH:

From Treating Illness to Predicting It

Dubai's health sector is deploying AI to detect disease earlier, run hospitals smarter and shift the whole system from reactive care to prevention



A reflection by H.E. Dr. Alawi Alsheikh-Ali, Director General, Dubai Health Authority

> *"The true value of artificial intelligence is not measured by how widely it is used, but by the tangible results and lasting positive impact it has on people's lives.*

Nowhere are the stakes of artificial intelligence higher, or the rewards more human, than in health. A model that flags a patient's risk earlier is not an efficiency statistic; it is a life redirected. Guided by our leadership, Dubai's digital transformation ambitions and the UAE National Strategy for Artificial Intelligence 2031, our vision is clear: to position Dubai as a global pioneer in disciplined, responsible and human-centred AI-driven healthcare.

The Dubai Health Authority has moved decisively from exploring AI to embedding it as a strategic capability across the healthcare ecosystem. Our first priority was the foundations that let AI scale safely: strengthened governance and regulatory frameworks, secure and interoperable data infrastructure, and a workforce equipped with AI and digital competencies. On those foundations we prioritized initiatives that change outcomes directly. Advanced analytics now support the prevention and early detection of chronic disease, with our diabetes program showing a 42 per cent improvement in results among at-risk participants. In our contact centres, the AI-enabled Customer Risk Radar identifies dissatisfaction before it hardens, deflecting up to 93 per cent of negative interactions.

The lesson of the journey

The most important lesson of the past three years is that successful AI adoption is not simply a technology initiative; it is an organizational transformation. Data quality, governance, workforce readiness and ecosystem collaboration matter as much as algorithms. Where AI is introduced with clear governance and aligned with real operational needs, it lifts productivity, strengthens compliance and improves the experience of almost every department. Sustaining that trust, through ethical and responsible AI practice, is continuous work, and we treat it as such.

The next phase is already defined. DHA will publish the AI Strategy for Healthcare Sector in Dubai, a three-year roadmap that moves the sector beyond adopting AI to advancing disciplined, intelligent health, through sector-wide governance frameworks, workforce readiness and AI-enabled services. Artificial intelligence is among the most transformative opportunities in modern healthcare. Dubai's intention is to show how it is done responsibly, at scale, and in the service of healthier communities

In this sector

Dubai Health Authority p. 77
Sets the strategic direction and governance of Dubai's healthcare sector, overseeing and regulating both public and private healthcare providers. It also leverages AI to improve population health outcomes, enhance service delivery and experience, and drive the responsible adoption and integration of AI across the healthcare ecosystem.

Dubai Health p. 78
Brings AI to the bedside and the medical record as the first integrated academic health system in Dubai.

REFERENCE
[1] Dubai Health Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.
[2] Dubai Health, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



4a

DUBAI HEALTH AUTHORITY:

Prevention, Predicted

Machine learning that catches diabetes before it develops and an AI radar that resolves dissatisfaction before it escalates: the Dubai Health Authority is applying AI where it changes outcomes



"Our vision is to shape a future-ready healthcare sector in Dubai, harnessing artificial intelligence to anticipate health risks, strengthen prevention and enhance the healthcare experience, while keeping people at the heart of care."

H.E. Dr. Alawi Alsheikh-Ali
Director General,
Dubai Health Authority

AI RISK RADAR:
CUSTOMER HAPPINESS, SINCE EARLY 2025

AT 6 MONTHS	NOW
-80%	-93%
negative calls	negative calls

97%
quality and compliance among staff

The Dubai Health Authority's AI portfolio is built around a simple conviction: the best healthcare intervention is the one that arrives before the illness. As the steward of Dubai's health sector, DHA has paired that conviction with the enablers that let AI scale safely, strong governance and regulatory frameworks, secure and interoperable data infrastructure, and a workforce trained in digital and AI competencies, and directed the technology at the sector's highest-value problems.

The clearest proof is in chronic disease. Since mid-2024, a machine learning solution in public health has identified individuals at risk of developing diabetes before the condition takes hold, operating at 78 per cent predictive accuracy. Participants flagged by the model receive preventive guidance early, and the results are striking: a 42 per cent improvement in outcomes among at-risk participants, with 21 per cent improving so far that they moved from the pre-diabetes stage back to normal blood sugar levels. Beyond individual lives changed, the program is projected to reduce the financial burden on the healthcare system by 10 per cent over the long term, and it is designed to scale to the whole Dubai population.

The same anticipatory logic runs through the service experience. In the Customer Happiness contact centre, the AI Risk Radar transcribes and analyses voice interactions in real time, extracting indicators that predict potential dissatisfaction before a call goes wrong. Within six months of its early-2025 deployment, calls classified as negative fell by 80 per cent, a figure that now stands at 93 per cent. Every call is monitored, quality and compliance among contact centre staff have risen to 97 per cent, and the platform absorbs peak-season demand without increasing operating costs.

Both deployments illustrate the DHA method: AI introduced with clear governance, aligned to a real operational need, and measured relentlessly. It is an approach the authority now intends to extend across the sector through the forthcoming AI Strategy for Healthcare Sector in Dubai, a three-year roadmap spanning sector-wide AI governance, workforce readiness, AI-enabled internal and external services, and an Intelligent Health Hub.

As a regulator and steward of population health, DHA is uniquely positioned to shape the future of Dubai's healthcare sector by harnessing AI to strengthen governance, empower healthcare professionals, accelerate innovation, and drive the shift from reactive treatment to proactive, predictive, and personalized care.

REFERENCE
[1] Dubai Health Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



4b

DUBAI HEALTH:

Intelligence at the Bedside

Inside the first integrated academic health system in Dubai, AI now guards patient safety, predicts deterioration and gives clinicians their time back

"Artificial intelligence is not viewed as a standalone technology initiative, but as a foundational capability embedded across our missions of Care, Learning, Discovery, and Giving."



Dr. Amer Sharif
Chief Executive Officer,
Dubai Health

Dubai Health, the first integrated academic health system in Dubai, has embedded artificial intelligence where healthcare is practised: inside Salama, Dubai Health's Epic-based electronic health record (EHR) that clinicians use throughout their day. The EHR has been building layers of intelligence for nearly a decade and today supports a spectrum of capabilities, from rule-based safety logic to generative AI.

The foundation layer has been supporting patient safety since 2017. Best Practice Advisory rules evaluate patient data in real time during ordering, prescribing and documentation, providing alerts for drug-to-drug interactions, prophylaxis reminders and immunization schedules.

Building on this foundation, Epic's predictive clinical risk models, deployed progressively since 2018, apply machine learning to forecast sepsis, patient deterioration, ICU readmission and other risks. Early insights are surfaced on clinical dashboards, helping care teams identify opportunities to act sooner and support patients at the right time.

AI is also giving clinicians more time to focus on what matters most: their patients. Fluency Direct speech recognition, integrated with the Salama EHR since 2023, enables physicians to dictate directly into the medical record with high transcription accuracy. Its documentation intelligence engine provides real-time feedback that supports coding completeness and clinical documentation quality. In November 2025, Dubai Health expanded these capabilities by activating Epic's generative AI tools, including the recent activation of the ambient clinical documentation that listens to the patient-physician conversation and generates the required clinical documentation. Large language models can draft clinical summaries and patient responses, assist coding and answer conversational data queries, with clinician validation before information enters the record. Together, these tools are helping streamline documentation and create more time for patient care.

The portfolio reflects a clear and deliberate approach. Dubai Health's experience shows that sustained AI transformation is supported by a balanced portfolio across clinical care, operations, workforce enablement and patient engagement, underpinned by organizational readiness, engaged clinical leadership, thoughtful change management and workflow integration. Working with partners such as Epic, Dubai Health has made measurement and transparency central to building institutional trust.

The direction from here is clear: responsibly scaling AI across Dubai Health under rigorous clinical oversight and ethical governance, supporting increasingly predictive and preventive models of care.

For Dubai Health, the true value of AI lies not in the technology itself, but in what it enables: empowering clinicians, enhancing the patient experience and advancing the health of the communities it serves.

REFERENCE
[1] Dubai Health, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.
[2] Dubai Health, "CEO message: AI as a foundational capability in an integrated academic health system," State of AI Report 2026 entity submission (supplementary), Dubai, UAE, 2026.

5

SECURITY & SAFETY:

The Shield Becomes Intelligent



A reflection by H.E. Lieutenant General Abdullah Khalifa Al Marri, Commander-in-Chief, Dubai Police, opening the Security & Safety cluster

Everything Dubai builds in the AI era rests on a single foundation: the confidence of the people who live in, visit, and invest in this city. Security is where that confidence is earned. Over the past three years, we have led a deliberate transformation within Dubai Police, elevating artificial intelligence from targeted enablement in 2023 to a fully institutionalised operational capability by 2025. What began in selected services is now an integrated, enterprise-wide ecosystem embedded across operations, cybersecurity, investigations, and service delivery. Our vision is a fully cognitive police force, intelligent, predictive, and digitally unified, where AI serves as a force multiplier for human judgment, deployed responsibly, securely, and at scale.

Scale under pressure

The environment we protect has never been more demanding. Attempted cyber breaches against our systems rose from approximately 10 million in 2021 to more than 74 million in 2025. Our answer was not to slow down; it was to govern harder. Dubai Police became one of the first government entities in the world to achieve ISO/IEC 42001:2023 certification for AI management, and every deployment is aligned with Dubai's AI Ethics Principles to ensure transparency, human oversight, and accountability. The results speak for themselves: zero successful cyber breaches, a security operations response time under three minutes, and interoperability across 87 partner entities handling 301 million transactions since 2021, validated by a community happiness rate of 99.99 per cent.

> *"Technology alone does not transform institutions; culture and governance do. Governance preceded every deployment, and the results speak clearly.*

The lesson we carry forward

Our journey reinforced a defining lesson: scaling AI is a matter of culture and governance as much as technology. Looking ahead, we are advancing toward agentic AI, autonomous systems, and next-generation computational capabilities. My commitment is unwavering: Dubai Police, together with our partners across this sector, will continue to lead globally in ethical, secure, and scalable AI, ensuring that intelligent empowerment remains a public trust we uphold for Dubai and the United Arab Emirates.

In this sector

Dubai Police p. 80
Runs AI systems now live across the city, from unified smart surveillance to AI-enabled patrols and autonomous aerial response.

General Command of Dubai Civil Defence p. 81
Replaces reaction with prediction through a first-of-its-kind AI fire-readiness program built by national talent on seven years of incident data.

Dubai Electronic Security Center pp. 82 to 83
Sets out the policy architecture and international engagement through which Dubai secures the AI era.

REFERENCE

[1] Dubai Police, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.

[2] General Command of Dubai Civil Defence, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.

[3] Dubai Electronic Security Center, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



5a

DUBAI POLICE:

The Cognitive Force

From city-wide surveillance to autonomous aerial response, Dubai Police has turned AI into everyday operational capability

> *"Our strategic vision is clear: to build a fully cognitive police force, intelligent, predictive, and digitally unified, where AI serves as a force multiplier for human judgment, deployed responsibly, securely, and at scale."*



H.E. Lieutenant General Abdullah Khalifa Al Marri
Commander-in-Chief, Dubai Police

Few institutions anywhere have moved AI from experiment to infrastructure as decisively as Dubai Police. In three years, guided by the three strategic directions of its Digital Transformation Strategy 2024 to 2033, the force has built an enterprise-wide AI ecosystem spanning operations, cybersecurity, investigations, and service delivery, and it now anchors the security of a city of more than 3.6 million residents and millions of annual visitors.

Live across the city

Three production platforms define the force's operational edge. OYOON, the world's first surveillance platform to consolidate all security system features in a single AI-powered environment, unifies facial recognition, number-plate recognition, and video analytics across government, semi-government, and private partner entities under an ongoing five-year city-wide expansion. GHIATH, the AI smart patrol system embedded in every Dubai Police patrol vehicle, combines computer vision with AI route optimisation; it has cut the error rate in receiving incident reports from 52 per cent to 10 per cent, holds 12 registered intellectual properties, and has been benchmarked by 28 entities across the UAE and internationally, alongside a 99.6 per cent reported feeling of safety. Drone Box, a city-wide strategic initiative launched by HH. Sheikh Mohammed bin Rashid Al Maktoum and the first BVLOS-licensed unmanned aerial system in the Middle East and North Africa, deploys AI-guided drones from fixed stations to reach incidents ahead of ground units, reducing emergency response times and fuel consumption while cutting unit costs to less than a tenth of helicopter deployment. Each platform is trademark- or IP-protected, benchmarked internationally, and architected for replication across the UAE and the GCC.

Governed to scale

Capability of this order demands institutional discipline. Dubai Police was among the first government entities worldwide to achieve ISO/IEC 42001:2023 certification for AI management, runs its systems on infrastructure delivering 99.998 per cent data-centre uptime, and has recorded zero successful cyber breaches even as attempted attacks rose above 74 million in 2025. Workforce readiness programs cut phishing susceptibility by 75 per cent, and integration with 87 partner entities has processed 301 million transactions since 2021, earning a community happiness rate of 99.99 per cent and external customer satisfaction of 93.8 per cent in 2024. Sustained investment in workforce readiness underpins it all, reflecting the force's defining lesson that culture and governance, not technology alone, transform institutions.

Dubai Police is now advancing toward agentic AI, autonomous systems, and next-generation computational capability, extending its position as the reference point for intelligent, trusted policing worldwide.

REFERENCE

[1] Dubai Police, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.

5b

GENERAL COMMAND OF DUBAI CIVIL DEFENCE:

Fighting Fires Before They Start

A first-of-its-kind AI fire-readiness program, built in-house by national talent, is turning civil defence from response into prediction

"Our strategy rests on a decisive shift from traditional models to proactive, predictive, data-driven civil defence, deploying intelligent technologies in prevention, response, and decision- making to achieve the highest levels of protection for lives and property."



H.E. Lieutenant General Expert Rashid Thani Al Matrooshi
Commander-in-Chief, General Command of Dubai Civil Defence

Dubai Civil Defence has adopted an integrated artificial intelligence strategy that applies smart technologies across prevention, response, decision-making, and institutional and community readiness. Its centrepiece is the Civil Defence Readiness Program, recognised as the first of its kind globally in the field of firefighting and rescue, registered as intellectual property, and developed entirely in-house by highly skilled national talent, a demonstration of Emirati capability leading government innovation and a working expression of the UAE AI Strategy 2031.

Seven years of data, one predictive engine

The program applies AI to seven years of incident big data to identify fire patterns and their causes, detect high-risk behaviours, and trigger proactive awareness messaging tailored to language, culture, nationality, and occasion. More than four million AI-generated awareness messages have been delivered, supported by a smart virtual officer able to interact in more than 170 languages, and campaigns that once took a month now reach more than a million beneficiaries in a single day. The results are measurable: fires per 100,000 population fell 27 per cent, from 10.49 in 2021 to 7.66 in 2024.

A city mapped for risk

A dynamic AI heat map fuses more than 100,000 records from over 20 official entities, covering buildings, warehouses, hazardous materials, and critical infrastructure, to direct inspection teams precisely toward the highest-risk sites. In the field, AI-equipped drones analyse smoke colour and fire direction in real time to guide incident commanders, and conduct preventive surveys of warehouses and high-risk zones, detecting unsafe storage and obstacles blocking fire-appliance access and converting findings into enforceable, trackable violations. Computer- vision systems continuously measure fire-station response times from the moment vehicles depart to the moment they arrive on scene, while parallel analytics in customer service centres track waiting times and peak loads to sharpen frontline performance.

Exporting readiness

Through its AI-powered One Billion Readiness initiative, the General Command is now extending this model globally, using multilingual smart avatars and adaptive training content to raise the preparedness of one billion people worldwide. Dubai's civil defence has moved decisively from reacting to fires to predicting them, and it intends to carry the world with it.

REFERENCE [1] General Command of Dubai Civil Defence, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.

5c

DUBAI ELECTRONIC SECURITY CENTER:

Securing the AI Era

Through binding policy architecture at home and thought leadership on the world stage, the Dubai Electronic Security Center is making trust the operating condition of Dubai's AI decade



"AI can only reach its full potential when innovation is built on security and trust. At DESC, we remain committed to embedding security throughout the AI lifecycle, strengthening Dubai's digital foundations, and contributing to global efforts to shape a secure and trusted AI future."

H.E. Yousuf Al Shaibani
Chief Executive,
Dubai Electronic Security Center



A city that runs its services, infrastructure, and economy on artificial intelligence has made a wager on trust. Someone must underwrite that wager. In Dubai, that institution is the Dubai Electronic Security Center (DESC), the emirate's electronic security authority and this report's security partner. While the entities profiled across this chapter show what AI can do, DESC's mandate is the harder, quieter discipline beneath them all: ensuring that every model, dataset, and deployment remains worthy of the confidence Dubai's residents place in their government. AI unlocks extraordinary progress, but it also introduces vulnerabilities of a new kind, from data leakage to manipulated models and malicious exploitation of the systems themselves. DESC's answer operates on two fronts at once: a binding policy architecture at home, and active engagement in the international forums where global cyber defence practice is shaped.

A policy architecture for the AI lifecycle

DESC's AI Security Policy is the emirate's foundational instrument for secure AI adoption, and its design reflects a distinctive insight: AI security cannot be inspected in at the end; it must be governed across the entire lifecycle. The policy follows AI systems through design, development, deployment, monitoring, and disposal, and assigns clear, differentiated obligations to each stakeholder in that chain. AI Providers, whether global model developers or in-house government teams, face security requirements from threat and vulnerability management through secure design and disposal. AI Consumers, the government and critical non-government entities that deploy AI, carry governance obligations spanning senior management, procurement, legal, and machine-learning operations. Even end-users are within scope, with defined requirements for access management and the use of publicly accessible AI systems.

Two features give the policy its force. First, it is mandatory for all AI Providers offering services to Dubai government and critical non-government entities: security is a condition of entry to the government market, not a preference. Second, it anchors sovereignty in plain terms: throughout the AI lifecycle, data must remain within the borders of the UAE, at rest or in motion, backed up or recovered. Aligned with Dubai's Information Security Regulation and data protection laws, and supported by detailed mappings of AI-specific threats and vulnerabilities, the policy gives every entity in this report a common security language, and it insists on continuous testing, evaluation, verification, and validation rather than one-time certification.



Empowering defenders on the world stage

WEF EMPOWERING DEFENDERS 2026

~USD 1.9M
reduced average breach cost

~80 days
shorter breach lifecycle

or organizations that use AI extensively in security

DESC's influence does not stop at the city limits. The Center is a contributing partner to the World Economic Forum's Cyber Frontiers initiative, which convened experts from 84 organizations across 15 industries and whose May 2026 report, Empowering Defenders: AI for Cybersecurity, provides organizations worldwide with actionable guidance on deploying AI for cyber defence. The collaboration addresses a decisive shift in the threat landscape: the Forum's Global Cybersecurity Outlook 2026 finds that 94 per cent of respondents expect AI to be the most significant driver of change in cybersecurity, and that 77 per cent of organizations already use AI in their cybersecurity operations. The stakes are equally concrete: 88 per cent of security teams report time savings from AI-driven automation, and organizations that leverage AI extensively in security have reduced average breach costs by USD 1.9 million and shortened breach lifecycles by approximately 80 days.

Within the report, DESC's contribution illustrates the governance function itself, showing how AI can assess system configurations, audit trails, and control implementations against regulatory frameworks, validate whether security policies are correctly implemented across on-premise, cloud, and hybrid environments, and generate audit-ready documentation with traceable evidence. It is a fitting role: the emirate that made AI security governance mandatory is now demonstrating to the world how AI can strengthen governance in return. The initiative's forward agenda, including oversight and guardrails for agentic AI, keeps DESC engaged in shaping the next generation of international guidance on AI for cyber defence.

Trust as infrastructure

> *"Dubai is not merely consuming the global playbook for AI security. It is helping to write it.*

The significance of this dual posture, binding policy at home and shared doctrine abroad, is easy to understate. Attackers now operate at machine speed, generating malware, exploiting vulnerabilities, and scaling attacks in minutes. A city cannot answer that threat with enthusiasm alone; it answers with architecture. DESC's AI Security Policy means that AI systems serving Dubai's government and critical non-government entities must be designed, deployed, and monitored against a common standard, and its international engagement means that standard evolves with the global state of the art rather than behind it.

Scaling with trust is this report's theme, and DESC is where the phrase becomes literal. In Dubai's AI era, security is not a constraint on ambition. It is the infrastructure that makes ambition safe to build on.

REFERENCE
[1] Dubai Electronic Security Center, "Artificial Intelligence (AI) Security Policy, Version 1," Dubai, UAE, 2024.
[2] World Economic Forum, "Empowering Defenders: AI for Cybersecurity," in collaboration with KPMG, Geneva, Switzerland, 2026.
[3] World Economic Forum, "Global Cybersecurity Outlook 2026," Geneva, Switzerland, 2026.

6

ECONOMIC GROWTH:

The Intelligent Economy



A reflection by H.E. Dr. Abdulla Mohammed Busenad Director General, Dubai Customs, opening the Economic Growth cluster

Dubai's economy has always been built on movement: goods across borders, visitors through terminals, capital toward opportunity. The question of our era is whether that movement can keep accelerating while becoming more secure, more transparent, and simpler for the people behind it. Over the past few years, Dubai Customs has answered with a clear and planned vision: to establish itself as a global leader in AI-powered trade facilitation, enhancing efficiency, security, and customer satisfaction. We recognised early that artificial intelligence would not simply optimise isolated processes; it would fundamentally reshape how we manage risk, assist traders, and safeguard revenue in an increasingly complex global trade environment. Our ambition has been to embed AI in the core of Customs operations, so that technology directly advances our mandate of seamless trade facilitation and comprehensive compliance.

From friction to flow

That conviction now runs through the core of our operations, supported by deliberate, multi-faceted strategic directions: strengthening data-driven decision-making, expanding predictive analytics, and investing substantially in the capabilities of our workforce. Our AI Munasiq platform, the world's first AI-powered goods classification engine of its kind, has compressed a task that once took a full day into seconds, serving users across 134 countries at a sustained 100 per cent service level and generating AED 34 million in operational cost efficiencies. Predictive analytics strengthen customs risk management, AI-enabled automation advances the government's Zero Bureaucracy agenda, and a comprehensive AI governance and ethical framework ensures that transparency, security, and accountability grow in step with capability. The lesson we draw is one this whole sector shares: intelligence, applied with discipline, converts friction into flow, and flow into competitiveness.

> *"Artificial intelligence would not simply optimise isolated processes; it would fundamentally reshape how we manage risk, assist traders, and safeguard revenue.*

The road ahead

Our own next chapter is already under way, with multi-agent AI orchestration being developed to enable coordinated, explainable decision-making across declaration workflows, from tariff validation and valuation to inspection recommendations and post-clearance audit. The destination is a smarter, more resilient, future-ready trade and visitor economy, one that sets the benchmark for how open economies grow: faster because they are trusted, and trusted because they are intelligently governed.

In this sector

Dubai Customs p. 85
Protects and accelerates the trade that anchors Dubai's position as a global logistics hub.

Dubai Sports Council p. 86
Turns AI toward the visitor economy and community life: a personalised sports itinerary before a traveller lands, public feedback heard in 19 languages, and event permits cleared in minutes.

Together they trace a single arc, an economy in which every journey, whether of a container or a family, is anticipated, assisted, and safeguarded by AI.

REFERENCE
[1] Dubai Customs, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.
[2] Dubai Sports Council, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



6a

DUBAI CUSTOMS:

Trade at the Speed of Thought

The world's first AI-powered goods classification platform has made Dubai Customs the global reference for intelligent trade facilitation

"Our ambition has been to embed AI in the core of Customs operations, ensuring that technology directly advances our mandate of seamless trade facilitation and comprehensive compliance."



H.E. Dr. Abdulla Mohammed Busenad
Director General,
Dubai Customs

Every shipment that enters Dubai must be classified under the Harmonized System before it can move. For decades, that meant expert review measured in hours. Dubai Customs has measured it in seconds since late 2024, when AI Munasiq, the world's first AI-powered HS code classification platform, entered full production, built on generative AI, computer vision, and retrieval-augmented generation trained on customs data to deliver accurate, explainable results.

AI Munasiq: a global benchmark

Available as a mobile and web application secured through UAE Pass, AI Munasiq lets traders, brokers, and businesses identify the correct tariff code through text, voice, or image input, in Arabic or English, together with duty rates, restrictions, and trade agreement benefits. Classification time has fallen from up to 24 hours to five to ten seconds. The platform operates at a sustained 100 per cent service level with 99.55 per cent uptime, its adoption rate has climbed from 71 per cent to 96 per cent with 2,300 active users across 134 countries, and it has generated AED 34 million in operational cost efficiencies while holding a 91 per cent customer happiness score and a 94.2 per cent service quality rating. Its AI models classify at 88 per cent accuracy, with the remainder refined through a human-in-the-loop reinforcement cycle that keeps precision improving in production. The Global Innovation Institute has endorsed the platform at its highest tier of value-creation innovation.

Intelligence across the declaration lifecycle

AI Munasiq is the flagship of a wider operational portfolio spanning eight strategic domains, from trade intelligence and predictive risk to inspection optimisation and customer service automation. The Customer Insights Platform, live since early 2026, applies machine-learning segmentation to give Dubai Customs a 360-degree view of its customers, enabling real-time compliance checks and saving 116 hours per loyalty enrolment request. A Predictive Risk Engine is strengthening customs risk assessment with advanced AI models, while an AI-enabled feedback analysis capability converts multi-channel customer sentiment into structured service improvement plans. A comprehensive AI governance and ethical framework, emphasising transparency, security, and accountability, governs the whole, and a hybrid infrastructure strategy pairing high-performance on-premise environments with scalable cloud capability keeps it resilient.

With multi-agent AI orchestration under development to coordinate tariff validation, valuation, inspection recommendations, and post-clearance audit, Dubai Customs is defining what the intelligent customs administration of the next decade looks like, from Dubai outward.

REFERENCE

[1] Dubai Customs, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



6b

DUBAI SPORTS COUNCIL:

The Most Intelligent Game in Town

From permit desks to personalised itineraries, Dubai Sports Council is building the region's most intelligent sports authority

"Our vision is to make DSC the most intelligent sports authority in the region, one where every service, every decision, and every citizen interaction is informed by data, powered by AI, and anchored in the values of Dubai's national AI strategy."



H.E. Saeed Hareb
Secretary General,
Dubai Sports Council

Dubai Sports Council (DSC) has moved from its first AI chatbot in 2023 to a portfolio of production systems that now touch every part of its mandate, from executive decision-making to the visitor experience. Each system began with a specific operational pain point, a manual process, a language barrier, or a data gap, and each is delivered against a simple test: AI succeeds when it solves real problems for real people rather than when it chases technology for its own sake.

Intelligence in daily operation

DSC leadership now interrogates the Council's entire sports database in plain Arabic or English through DSC AI Analytics, a natural-language platform that returns charts and executive summaries in seconds at over 90 per cent query accuracy, while the six-module AI Analytics Portal generates board-level intelligence reports automatically, saving more than 400 staff hours a year. In licensing, an AI permit recommendations engine pre-screens every incoming sports event application, cutting review time per permit by 40 per cent and catching compliance gaps that manual review previously missed. The multilingual Smart Survey System gathers public feedback in 19 languages, including voice, with AI translation and sentiment analysis opening DSC's ears to Dubai's more than 200 nationalities for the first time. A 24/7 WhatsApp channel handles thousands of public enquiries, and a live integration with the Department of Economy and Tourism's trade licence registry validates licences automatically at the point of event registration.

AI for the visitor economy

The public face of this transformation is the AI Sports Planner, available through Riyada, DSC's AI assistant on dubaisc.ae and on WhatsApp, live since late 2025. Choosing "Plan a sports trip", visitors and residents enter their dates and location and receive a fully personalised sports itinerary, drawn from DSC's live events and establishments database and delivered in seconds, covering nearby sports events, fitness centres, and healthy dining, and putting Dubai's sporting life in every traveller's pocket before they arrive.

Sovereign by design

DSC was among the first Dubai government entities to begin onboarding onto Digital Dubai's sovereign LLMaaS platform, bringing its AI workloads onto UAE-based infrastructure in full compliance with national data residency requirements.

With 54 AI use cases mapped across 10 strategic domains under its ICT Strategic Plan 2025–2030, DSC is now building toward its declared ambition: to make Dubai the premier Sports AI Hub for the MENA region by 2030.

REFERENCE

[1] Dubai Sports Council, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.

7

KNOWLEDGE & HUMAN DEVELOPMENT:

Intelligence in Service of People



A reflection by H.E. Hessa Bint Essa Buhumaid, Director General, Community Development Authority (CDA)

Every chapter of this report describes systems: platforms, models, engines. This one begins somewhere else, with a senior citizen whose nutrition is followed with new attentiveness, a family whose case is seen whole rather than in fragments, an artist whose work reaches its audience through services that simply work. Knowledge and human development is the sector where artificial intelligence meets its ultimate test, because here the beneficiary is not a process or a shipment but a person, often at a moment of need. If the preceding clusters showed how Dubai moves, powers, and protects itself with AI, this one asks the deeper question: what is all of it for?

Care that compounds

At the Community Development Authority, our mandate is the social fabric of Dubai itself: its families, its senior citizens, its people of determination, its most vulnerable. We have approached AI accordingly, not as an end but as a means of giving our social workers more time, more context, and more foresight. Intelligent tools now support nutrition care for our senior citizens, unify case management so that every family's story is understood in full, and relieve our teams of routine administrative burden so their energy goes where no machine can: into human judgement, empathy, and presence. The principle guiding every deployment is simple. Technology must serve care, never substitute for it. In social development, a system is only as good as the trust a family places in the person across the table, and everything we automate is chosen to protect that trust.

> *"The measure of intelligent government is not the sophistication of its systems but the wellbeing of the people they serve.*

The human decade

The value of AI will ultimately be judged not only by the intelligence of its systems, but by the strength and wellbeing of its communities: whether senior citizens enjoy healthier, more active lives, families become more resilient, people of determination find greater opportunities for inclusion, and empowered individuals flourish.

Achieving this balance requires more than technological investment. It calls for institutions and professionals who can combine the possibilities of AI with a deep understanding of human needs, social realities, and community aspirations. As technology advances, our responsibility to apply it wisely, inclusively, and compassionately grows with it. This is the future we seek: a Dubai where intelligence accelerates progress, while humanity continues to give that progress its purpose.

In this sector

Community Development Authority p. 88
Brings AI into the daily practice of social care.

Dubai Culture and Arts Authority p. 89
Applies AI to understand its audiences more deeply, to program culture around evidence of what communities actually seek, and to give Dubai's creatives secure, reliable digital services on which a creative economy can be built.

Culture and community development are twin guardians of the same asset: the human capital on which every other ambition in this report depends.

REFERENCE

[1] Community Development Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



7a

COMMUNITY DEVELOPMENT AUTHORITY:

AI for Community Wellbeing

Using intelligent tools to support professionals and strengthen services for individuals and families.

> *"At CDA, we embrace technology to advance our people's capabilities to understand and respond to community needs. By doing so, we empower individuals and families to lead a flourishing life."*



H.E. Hessa bint Essa Buhumaid
Director General, Community Development Authority

The Community Development Authority's approach to AI connects practical service needs with the expertise of social care professionals. Areas highlighted in the Authority's submission include nutritional support for senior citizens, unified case management, and assistance with everyday organisational tasks. Their shared purpose is to support more informed and responsive services while preserving the relationships on which effective social work depends.

Supporting senior citizens' wellbeing

The Thukher Nutrient application supports nutritional care within the Authority's senior citizens programmes. The submission describes intelligent tools intended to help professionals follow nutritional needs more consistently and give appropriate attention to individual circumstances. This application sits within a broader ambition to support senior citizens' quality of life and independence. Its value should be assessed through its contribution to professional practice and the experience of those receiving support, alongside the human relationships that sustain trust and belonging.

Understanding the context behind each case

CDA's Unified Case Management System brings relevant interactions and information into a more coherent view. AI assistance described in the submission supports tasks such as searching, drafting, and summarising information, helping professionals navigate records more efficiently.
A connected view of information can support continuity of service and more informed assessment. Professional judgement remains essential to understanding family dynamics, cultural context, and individual needs.
AI-generated material requires review by authorised staff, who remain responsible for assessments and decisions.

Developing professional capability

AI assistance with drafting, summarising, and organising information offers opportunities to reduce routine administrative demands. The intended benefit is to give staff more capacity for work that requires their expertise and direct engagement with the community.

Making effective use of these tools requires practical AI skills, critical assessment of outputs, and a clear understanding of appropriate use. These capabilities must develop alongside social knowledge and cultural awareness, so that improvements in efficiency translate into better support for people.

Responsible development and meaningful evaluation

The Authority is exploring opportunities to integrate models made available through Digital Dubai's large language model initiative into its services. Any integration should be assessed against service needs, information security requirements, and applicable government AI policies, with clear human oversight.

As these applications develop, evaluation should distinguish between intended benefits and demonstrated results. Useful measures include the quality and timeliness of support, administrative time saved, and the experience of individuals and families. Protecting sensitive information and maintaining accessible services remain central to building trust in their use.

REFERENCE

[1] Community Development Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



7b

DUBAI CULTURE AND ARTS AUTHORITY:

The Data Behind the Muse

Dubai Culture is building the intelligent foundations of a creative economy: data-informed programming and secure digital services for the culture sector

> *"Dubai Culture has approached artificial intelligence not as a standalone technology initiative, but as a strategic enabler of our broader vision and mission to preserve, promote, and advance culture, creativity, and knowledge through innovative public services."*



H.E. Hala Badri
Director General, Dubai Culture and Arts Authority

Culture would seem the last sector to run on algorithms. Dubai Culture and Arts Authority has proved the opposite: that a cultural institution serving millions of visitors, thousands of creatives, and a portfolio of heritage sites, libraries, and venues needs intelligence in its foundations precisely so that the art on top can flourish. Between 2023 and 2025 the Authority implemented more than ten AI applications across its key functions, achieving 92 per cent user satisfaction and 95 per cent user trust while generating savings of approximately AED 6 million.

Programming by evidence

At the heart of the portfolio is an AI analytics capability that gathers visitor information across the Authority's sites and applies machine learning to reveal how communities actually engage with cultural offerings. The result is cultural programming informed by evidence: understanding which experiences draw which audiences, and where the next investment of creative energy will matter most. For an authority whose mandate spans everything from heritage to contemporary art, replacing intuition with insight changes what it can offer the city.

Reliable rails for a creative economy

The rest of the portfolio works behind the curtain, keeping the sector's digital services trustworthy. Predictive AI continuously monitors the environmental conditions of core communications infrastructure, intervening before faults become outages for the public and creative communities that depend on the Authority's platforms. Cloud-based AI services accelerate how new digital services for the culture sector are built and delivered, while AI-driven risk-based access management and automated security operations protect the data of artists, institutions, and audiences around the clock. None of these systems hangs in a gallery; all of them determine whether a creative's permit, a library member's account, or a festival registration simply works.

People before platforms

The Authority paired its technology with more than 140 training courses for its workforce, reflecting a lesson from its own journey: transformation depends less on technology than on readiness across governance, culture, and skills, with trust designed into every solution from the beginning, especially where data, ethics, and user experience meet.

With its Digital Transformation Strategy 2026–2030 now launching an enterprise-wide intelligent transformation, Dubai Culture is set to prove that the world's most creative cities will also be its most intelligently run.

REFERENCE

[1] Dubai Culture and Arts Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026.



8

The Government-Wide AI Maturity Model

From pilot to production to optimization: where the centre of gravity of Dubai government AI now sits, and what distinguishes the entities furthest ahead

The twenty entity spotlights in this chapter reward two readings. Read singly, they are portraits of institutions applying AI to their own mandates. Read together, they yield something more useful: an empirical maturity model for government AI, drawn from the portfolio itself rather than imported from consultancy literature. Three stages recur across the submissions behind this chapter [1]. Entities pilot, then they produce, then they optimize. The stages are not a league table, and this report deliberately does not assign named entities to them. They are a lens for understanding where Dubai government stands as a whole, and what the institutions furthest along the curve do differently.

> *"The centre of gravity of Dubai government AI no longer sits in the laboratory. It sits in live operations, serving residents in real time.*

Three stages, one direction

The pilot stage is exploration under controlled conditions. A use case is scoped, data is assembled, feasibility is tested, and the organization learns what the technology can and cannot do for its mandate. Every entity in this chapter has passed through this stage, and most still maintain an exploratory frontier alongside their live systems; that is healthy, not a deficiency.

The production stage begins when a system leaves the laboratory and takes on real users, real transactions, and real accountability. It must now survive peak loads, edge cases, audits, and the expectations of residents who never volunteered to be testers. Production is where AI stops being a technology program and becomes a public service.

The optimization stage is quieter and rarer. Systems are instrumented, feedback loops are formalised, models are retrained on operational data, and the application estate is managed as a portfolio, with retirement as available a decision as expansion. At this stage the question changes from "does it work?" to "is it still the best way to deliver this outcome?"

The government-wide AI maturity model: three stages



Figure 3.8: The government-wide AI maturity model: three stages.

The question shifts from "does it work?" to "is it still the best way to deliver this outcome?"

The centre of gravity: production

On the evidence assembled here, the centre of gravity of Dubai government AI sits firmly in production. More than ninety production AI applications are documented across the twenty featured entities, spanning conversational services, predictive operations, computer vision in the field, and document intelligence [1]. These are live systems with measurable throughput and named owners. The alignment finding presented in Chapter 2, that 76.7 per cent of AI initiatives are mapped to entity or city strategy, is visible here at ground level: production systems attached to mandates, not experiments attached to enthusiasm.

What optimizing entities do differently

> *What separates optimizing entities is not the size of their portfolios but the discipline of their feedback loops.*

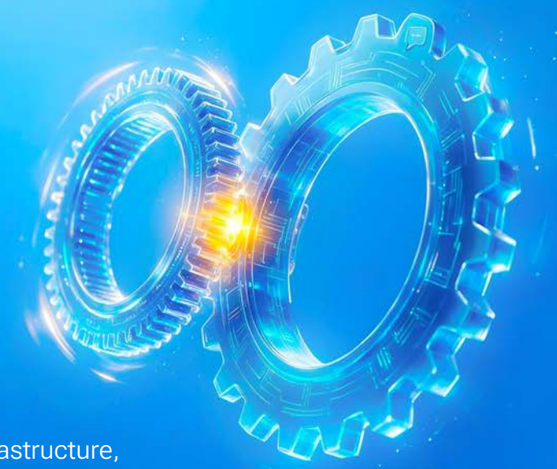
A smaller group of entities has moved beyond running systems to improving them continuously, and four traits distinguish this group. First, they measure, and they publish what they measure. The clearest illustration is DEWA's Rammas, which has handled more than 12.7 million user questions, returns hundreds of thousands of service hours annually, and delivers reported yearly savings of AED 12 to 14 million; a decade of iteration sits behind those figures. Second, they govern. Chapter 2 recorded that 80 per cent of entities report operating AI governance boards; in optimizing organizations those boards gate model changes and retraining, not only new procurements. Third, they build. Entities that maintain in-house engineering capability, a current that runs visibly through this chapter, shorten their improvement cycles from procurement timescales to sprint timescales. Fourth, they share, both consuming the shared AI services that just over half of entities already use and contributing patterns back to the rest of government.

A moving model

No entity occupies a single stage. Each runs a mixed portfolio: mature workhorses, newly productionised services, and an exploratory edge, and the healthiest portfolios in this chapter show all three layers at once. The model's value is therefore directional rather than diagnostic. It shows a government that has, in aggregate, crossed from proving AI to operating it, and that now faces the characteristic demands of the next stage: measurement, governance, and continuous improvement at scale. It also gives planners a common vocabulary. When entities describe their ambitions for the coming year, the language of the submissions is consistently the language of the third stage: instrumentation, retraining, consolidation, and reuse [1]. Those demands, and the trust obligations they carry, are the subject of the chapter that follows.

REFERENCE

[1] Digital Dubai Authority, "State of AI Report 2026 entity submissions analysis," Dubai, UAE, 2026.



9 Two Speeds, One System

Utilities and security run AI as industrial infrastructure, service sectors have made it the front door of government, and the constraints in between are shared

Comparing the seven sector clusters of this chapter is not an exercise in ranking; every sector shown here operates AI in production. It is an exercise in texture. Sectors have moved at different speeds, along different paths, because their mandates demand different things of the technology [1]. Read side by side, they reveal a city adopting AI in two registers at once, industrial and conversational, and converging on a common set of constraints.

INDUSTRIAL SCALE

Utilities and security

> *Utilities and security now operate AI at industrial scale. The service sectors have made conversational AI the front door of government.*

The deepest AI operations in Dubai government sit where the physical stakes are highest. In the utilities sector, DEWA documents thirty-three production applications, a portfolio that runs from solar generation forecasting through network operations to the Rammas assistant and its 12.7 million plus handled questions [1]. This is AI as infrastructure: embedded in the grid, measured continuously, and improved as a matter of routine. The security and safety cluster shows the same industrial character in a different medium. Dubai Police operates computer vision at city scale through OYOON, and Dubai Civil Defence has moved firefighting from response to anticipation with fire-pattern prediction built by its own engineers. What unites these sectors is density: many models, tightly integrated with operations, run by teams for whom AI is a core operational discipline rather than a digital service channel.

CITIZEN SCALE

The service sectors

The government services, health, economy, and knowledge clusters have concentrated their energy where residents feel it most directly: the interaction itself. Conversational and generative AI is the signature of these sectors. GDRFAD's Salama assistant, embedded in Dubai AI, guides residency journeys, MBRHE's retrieval-augmented assistant answers housing queries from verified policy, DOF's AILTIZAM reads and reconciles financial documentation, and Dubai Customs' Al Munasiq brings the same intelligence to trade declarations [1]. In health, the emphasis is shifting from access to anticipation, with chronic disease early detection at DHA a leading example. The service sectors demonstrate a distinct excellence: taking frontier generative technology and making it safe, multilingual, and dependable enough to stand between a government entity and the public.

BOTH REGISTERS

Mobility and urban

The mobility and urban cluster shows why the two registers are complementary rather than competing. RTA assesses road condition with LiDAR-based analysis, industrial observation in the physical world, while Dubai Municipality runs a broad portfolio of production applications across city services, and DCAA demonstrates the in-house engineering depth that the security and safety cluster also displays [1]. A sector responsible for both infrastructure and daily resident experience adopts in both registers at once, and its trajectory suggests where the other clusters converge.

Two registers, seven clusters: how Dubai's sectors adopt AI

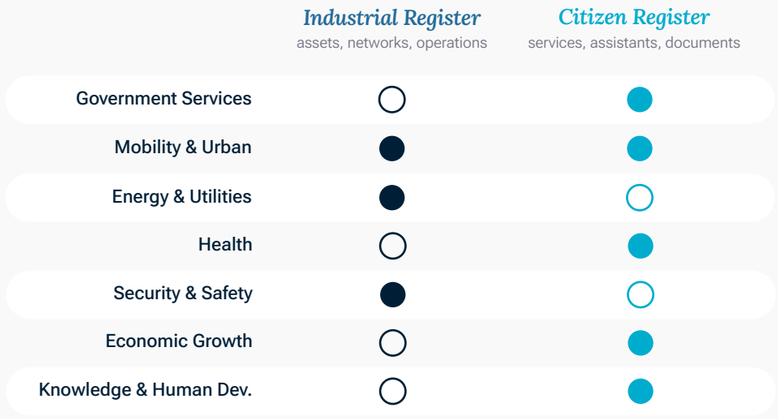


Figure 3.9:Two registers, seven clusters: how Dubai's sectors adopt AI. Filled circle marks the register where a cluster's production applications sit.

Shared bottlenecks, shared answers

The submissions are candid about what slows the portfolio down, and the striking finding is how uniform the constraints are across all seven clusters. Three recur government-wide [1]. Data readiness comes first: production AI consumes clean, governed, well-described data faster than legacy records practices produce it, and every sector reports effort flowing into data foundations. Specialist talent comes second: machine learning engineering, model operations, and AI risk skills remain scarce in every labour market on earth, and Dubai government competes for them like everyone else. Legacy integration comes third: the hardest engineering in most projects is not the model but the connection to systems of record designed decades before AI.

None of these constraints belongs to any single entity, and none is being solved entity by entity. Just over half of entities already consume shared AI services, the government-wide platforms described in Chapter 2 are turning common problems into common infrastructure, and two thirds of entities participate in AI research or benchmarking that circulates hard-won lessons [1]. The direction of travel is clear: a constraint met once, at city level, does not have to be met twenty times. The bottleneck analysis therefore ends in an unusual place for a comparison piece: the sectors differ most in what they have built, and least in what they must overcome. That shared ground is precisely what makes a shared response, and a single city-level maturity trajectory, possible.

> *"The bottlenecks are shared, and so, increasingly, are the solutions: common platforms, shared services, and pooled expertise.*

REFERENCE

[1] Digital Dubai Authority, "State of AI Report 2026 entity submissions analysis," Dubai, UAE, 2026.

10

What the Data Reveals

Five patterns recur across more than ninety production systems in twenty entities, and the strongest of them is convergence

> *"Five patterns recur across ninety-plus production systems, and none of them was mandated from the centre.*

Strip away the sector labels and read the more than ninety production applications documented in this chapter as a single dataset, and structure appears [1]. Five patterns recur with a regularity that individual spotlights cannot show. None was mandated from the centre; each emerged because separate entities, solving separate problems, kept reaching for the same answer. That independence is what gives the patterns their evidential weight. A directive produces uniformity by instruction; a portfolio produces it only when the underlying logic of the technology, and of government work itself, points the same way from every starting position.



PATTERN ONE

The conversational front door

Conversational and generative AI assistants are a prominent recurring pattern in the portfolio. DEWA's Rammas, with more than 12.7 million user questionshandled, is the longest-running example; GDRFAD's Salama and MBRHE's retrieval- augmented assistant show the same pattern in residency and housing [1]. The design logic is consistent: place a governed, multilingual conversational layer in front of complex services, ground it in verified content, and measure it relentlessly. The front door of Dubai government is increasingly a dialogue.

PATTERN TWO

From response to anticipation

A second prominent pattern is predictive and preventive operations. Dubai Civil Defence predicts fire patterns before incidents occur, DHA detects chronic disease earlier than symptoms would reveal it, and DEWA forecasts solar generation to balance the grid [1]. Different domains, identical shift: institutions that once excelled at responding are learning to act before the event. This is the pattern with the most headroom, because almost every mandate in government contains something worth anticipating.

"The strongest signal in the data is convergence: entities solving different problems keep arriving at the same architectures.

THE META-PATTERN: CONVERGENCE FROM ABOVE

80%
of entities report a governance board

2/3
participate in research or benchmarking

30
AI partnerships recorded in the 2025 policy self-assessment

PATTERN THREE

Vision in the field

Computer vision has left the laboratory and gone outdoors. Dubai Police's OYOON applies vision at city scale for public safety, while RTA assesses road condition with LiDAR-based analysis rather than manual inspection [1]. The common thread is the automation of observation: tasks that once required patrols and site visits are becoming continuous, consistent, and digital. For the entities involved, the gain is not only efficiency but coverage, because a city observed continuously is a city that can be maintained and protected before problems mature.

PATTERN FOUR

Document and case intelligence

A quieter pattern runs through the back office. DOF's AILTIZAM reads financial documentation, Dubai Customs' AI Munasiq interprets trade declarations, and CDA applies AI to social case management [1]. Government runs on documents and cases, and entities have converged on the same response: let machines do the reading, so professionals can do the judging. This class attracts fewer headlines than the conversational systems, yet it may carry the largest share of quiet productivity gain in the portfolio, because it compounds through every transaction an entity processes.

PATTERN FIVE

The in-house current

The portfolio is not simply procured. A strong engineering current runs through it: Dubai Civil Defence and DCAA built flagship systems with their own teams, and in-house capability appears across the submissions as a marker of the most mature operations [1]. It coexists with a deep ecosystem, with thirty-seven named company partnerships reported by ten entities, spanning firms from Microsoft and OpenAI to SAP, IBM, Epic, Capgemini, and SAS. The pattern is a deliberate blend: build the differentiating core, partner for the rest.

THE META-PATTERN

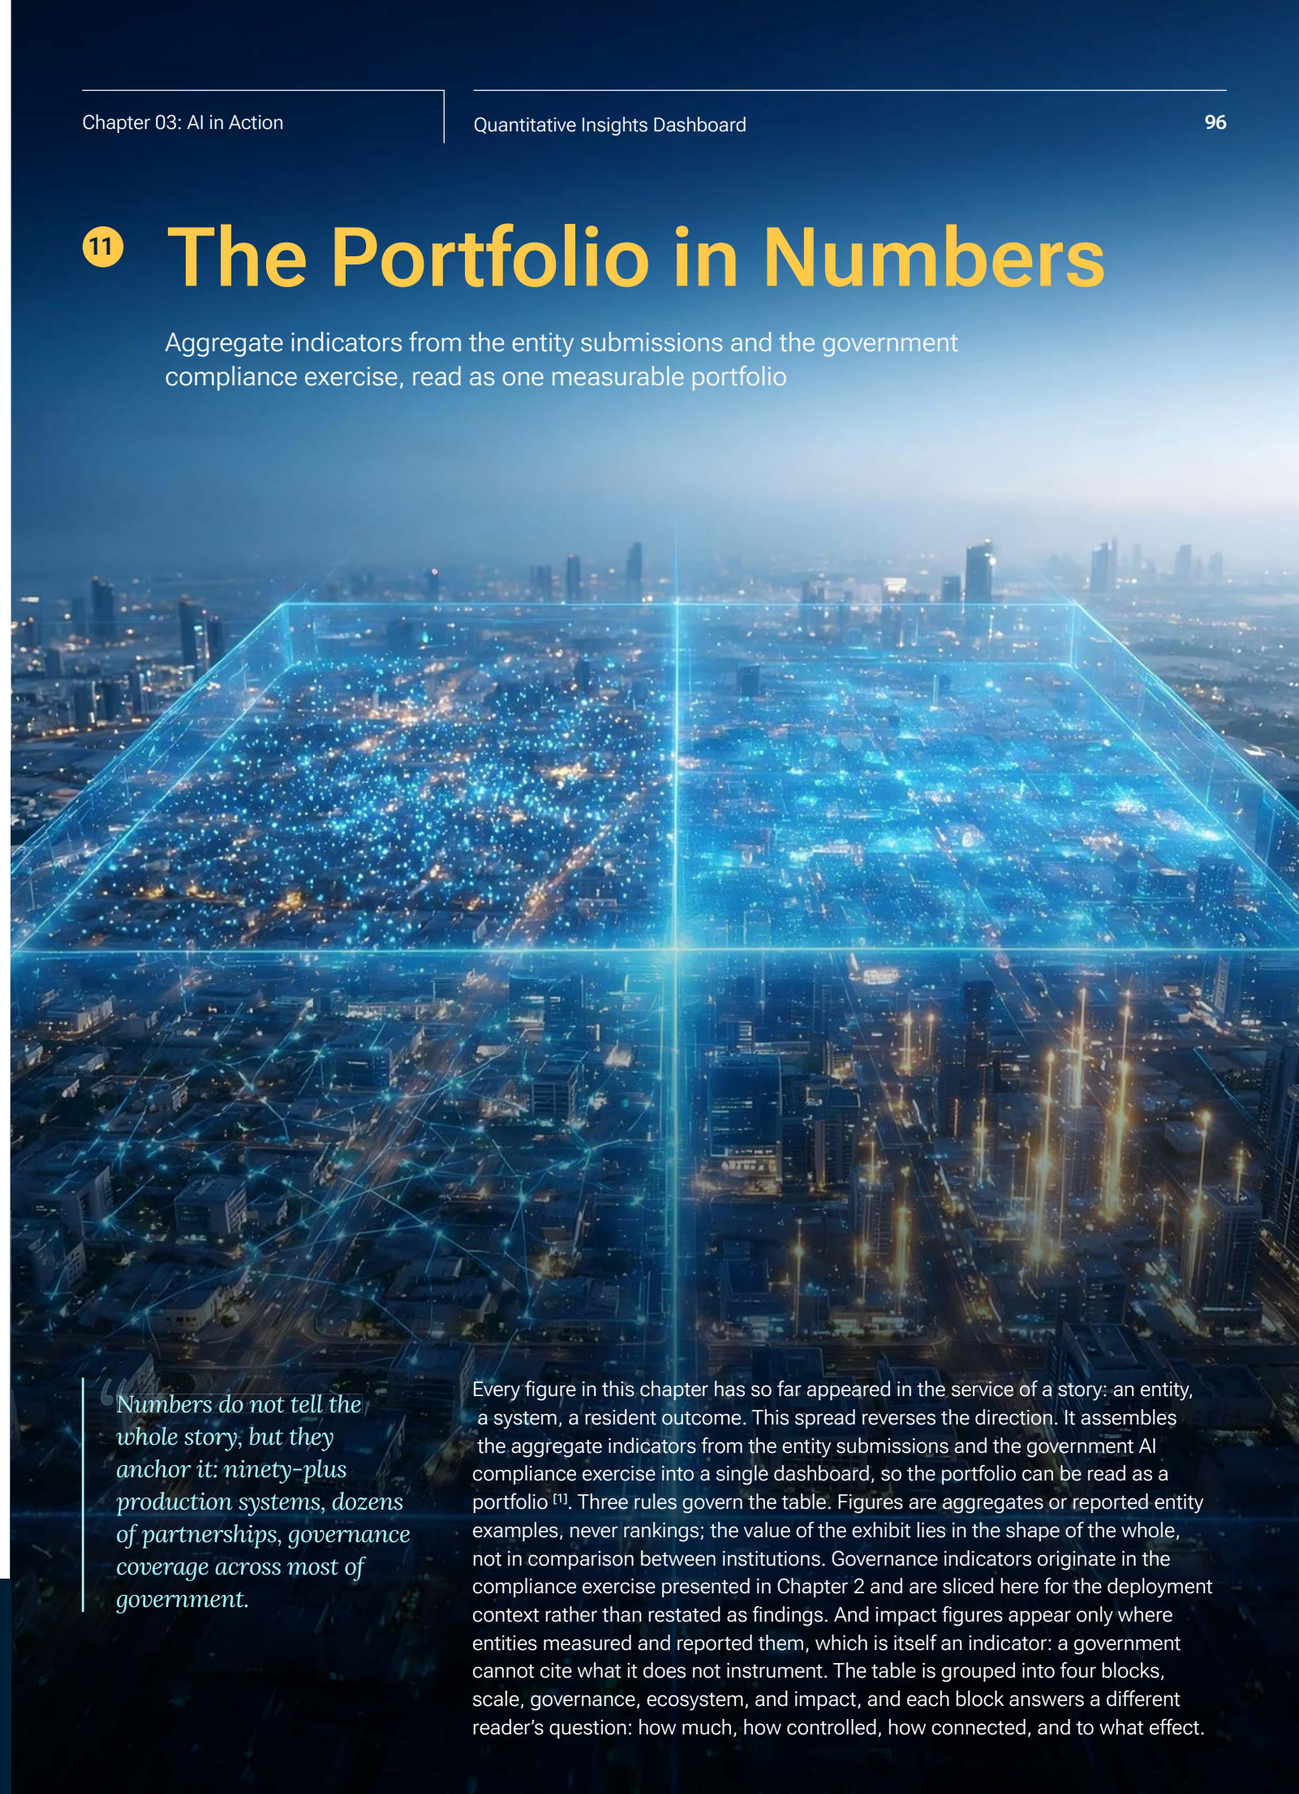
Convergence

The five patterns share one property: they emerged independently and converged anyway. Entities with different mandates, data estates, and histories arrived at the same architectural answers, which is what a maturing discipline looks like. Convergence also changes the economics of what comes next. When many entities run the same class of system, a lesson learned once, whether about grounding a conversational model, validating a prediction, or governing a vision feed, can be transferred at low cost, and shared services stop being a convenience and become a multiplier. The institutional aggregates presented in Chapter 2 make the same point from above: governance boards reported in 80 per cent of entities, two thirds participating in research or benchmarking, thirty AI partnerships recorded in the policy self-assessment [1]. Convergence from below, coordination from above. What the data reveals, in the end, is not twenty AI programs. It is one, distributed.

REFERENCE
[1] Digital Dubai Authority, "State of AI Report 2026 entity submissions analysis," Dubai, UAE, 2026.

11 The Portfolio in Numbers

Aggregate indicators from the entity submissions and the government compliance exercise, read as one measurable portfolio



"Numbers do not tell the whole story, but they anchor it: ninety-plus production systems, dozens of partnerships, governance coverage across most of government.

Every figure in this chapter has so far appeared in the service of a story: an entity, a system, a resident outcome. This spread reverses the direction. It assembles the aggregate indicators from the entity submissions and the government AI compliance exercise into a single dashboard, so the portfolio can be read as a portfolio [1]. Three rules govern the table. Figures are aggregates or reported entity examples, never rankings; the value of the exhibit lies in the shape of the whole, not in comparison between institutions. Governance indicators originate in the compliance exercise presented in Chapter 2 and are sliced here for the deployment context rather than restated as findings. And impact figures appear only where entities measured and reported them, which is itself an indicator: a government cannot cite what it does not instrument. The table is grouped into four blocks, scale, governance, ecosystem, and impact, and each block answers a different reader's question: how much, how controlled, how connected, and to what effect.

The portfolio in numbers: quantitative insights dashboard

PORTFOLIO SCALE

90+
production AI applications documented
across the 20 entities featured in this chapter

20
entities featured
organized in seven sector clusters

GOVERNANCE AND ALIGNMENT

80%
entities reporting AI governance boards
as reported in the 2025 policy self-assessment

76.7%
AI initiatives mapped to entity or city strategy
deployment follows mandate

53.3%
entities using shared AI services
common constraints, common infrastructure

66.7%
entities in AI research or benchmarking
lessons circulate across government

ECOSYSTEM ENGAGEMENT

30
AI partnerships recorded across government
from the government compliance exercise

37
named company partnerships
reported by 10 entities; including Microsoft, OpenAI, SAP, IBM

ADOPTION AND IMPACT

12.7M+
user questions handled
(DEWA Rammas)
cumulative, longest- running assistant

344K
service hours returned in 2025
(Rammas)
with AED 12–14M reported yearly savings

Figure 3.11: The portfolio in numbers: quantitative insights dashboard.

FOUR STATEMENTS

1. **Scale**
   An operating estate, not a collection of showcases.
2. **Discipline**
   The estate grew up inside institutions, not ahead of them.
3. **Openness**
   Built with the global ecosystem, not in isolation.
4. **Measurement**
   Come from systems instrumented from the start.

What the numbers say together

Read as one exhibit, the dashboard makes four statements. The first concerns scale: more than ninety production applications across twenty entities means AI in Dubai government is no longer a collection of showcases but an operating estate, broad enough that its behaviour is now a statistical matter rather than an anecdotal one [1]. The second concerns discipline: the governance block shows reported oversight boards in four fifths of entities and strategy mapping across three quarters of initiatives, so the estate has grown up inside institutional structures rather than ahead of them. The third concerns openness: thirty AI partnerships recorded in the policy self-assessment and thirty-seven named company relationships reported by ten entities in their submissions describe a portfolio built with the global technology ecosystem, not in isolation from it. The fourth concerns measurement itself. The strongest impact figures in the chapter, the 12.7 million questions handled by Rammas and the hundreds of thousands of service hours returned annually, come from systems that were instrumented from the start [1]. The lesson generalises: impact evidence is a design decision, made early, and the entities that made it now hold the strongest case for further investment of effort and attention.

From dashboard to discipline

A table of this kind also sets an institutional standard. Publishing aggregate indicators in a government report commits the government to being measured against them in the next edition, and that is the intent. The shared-services figure invites movement from just over half of entities towards all of them; the research participation figure invites the remaining third into the benchmarking community; the governance figure invites completion of coverage. The dashboard is therefore not only a record of the 2026 position but an instrument of the trajectory, a baseline against which Scaling with Trust can be tested in public.



Aggregate indicators turn twenty individual stories into one measurable portfolio.

What the numbers cannot say

A dashboard flattens, and honesty requires stating what this one omits. It does not capture quality of integration, the depth behind a deployed system, or the human redesign of work that determines whether a model changes outcomes. It reports adoption and impact only where entities measured them, so the impact block understates the whole. And it is a photograph of July 2026: several indicators, particularly shared-services uptake and predictive deployments, are moving quickly enough that the next edition should treat these values as a baseline [1].

The dashboard closes the evidence of this chapter in the appropriate register: measured, aggregate, and verifiable. Numbers of this kind are the precondition for the accountability the next chapter demands, because scrutiny without measurement is only opinion. The next spread asks what a government that can produce such a table owes its residents in return.

REFERENCE [1] Digital Dubai Authority, "State of AI Report 2026 entity submissions analysis," Dubai, UAE, 2026.



12

From Evidence to Scrutiny

Production scale makes trust questions concrete: a government that runs AI at this scale must also govern it at this scale

> *A government that runs AI at this scale must also govern it at this scale. The next chapter asks how.*

This chapter set out to answer a single question with evidence: is AI in Dubai government real? Across twenty entities, seven sectors, and more than ninety production applications, the answer is documented rather than asserted [1]. Systems answer millions of questions, forecast the grid, watch the roads, read the case files, and anticipate emergencies. The maturity analysis placed the centre of gravity in production; the pattern analysis showed independent entities converging on shared architectures; the dashboard rendered the whole as one measurable portfolio. The evidence phase of this report is complete. What follows from it is not celebration. It is obligation, and this closing spread explains why the report turns next to scrutiny.

Scale changes the question

When AI was experimental, trust was a principle to be affirmed. Now that it is operational, trust is a property to be engineered, and every pattern this chapter documented carries its own version of the question. Conversational systems standing between government and resident must answer accurately, in the right language, to the person entitled to ask. Predictive systems that act before the event must be right about the future often enough to justify acting on it, and inspectable when they are not. Vision systems that observe continuously must be governed as carefully as they are engineered. Document intelligence that reads at machine speed must preserve the human judgement it was built to serve. None of these is a hypothetical. Each is a live operating condition of systems described, positively and accurately, in the preceding forty pages [1].

> *Production is where trust stops being a principle and becomes a daily operational question.*

THE BRIDGE THREE FACTS

1. **Scale is real**
   This is a government operating AI, not piloting it.
2. **Convergence is real**
   Shared architectures mean risks and remedies are addressed once for the city, not twenty times
3. **Governance instinct is present**
   Visible in the aggregates and in the discipline of the most mature operators.



The questions production makes concrete

Production scale converts abstract AI ethics into specific institutional questions. Who is accountable when an automated decision is wrong, and how quickly can it be corrected? How is a model retrained, by whom, and under whose sign-off? Where does human oversight sit in a process that no longer has natural pauses? How does a government prove, rather than declare, that its systems are fair, secure, and explainable? These questions are not signs of weakness in the portfolio; they are the portfolio's consequence. They are also the five trade-offs of this report's analytical lens made operational: speed against safety, automation against oversight, innovation against regulation, performance against explainability, centralisation against autonomy. Chapter 3 showed those trade-offs being navigated in practice. Chapter 4 examines them directly.

Why the report stress-tests itself

It would be easy for a government publication to end here, on the high ground of evidence. This report deliberately does not. The chapter that follows is a stress test: an examination of risk, governance, assurance, and scrutiny conducted with the same rigour this chapter applied to capability. The reasoning is contained in the report's title. Scaling without trust is speed without steering; trust without scale is principle without proof. Dubai's position, documented across this chapter, is that the two must be built together, and the compliance foundations already in place, from governance boards to strategy mapping, show the scaffolding is real [1]. What remains is to test that scaffolding under load, openly and in public, which is what a report subtitled Scaling with Trust owes its readers.

The bridge

Readers should carry three facts across the chapter break. First, the scale is real: this is a government operating AI, not piloting it. Second, the convergence is real: entities are arriving at shared architectures and shared services, which means both risks and remedies can be addressed once, for the city, rather than twenty times. Third, the governance instinct is already present in the evidence, visible in the aggregates and in the discipline of the most mature operators. What Chapter 4 adds is pressure: the deliberate application of hard questions to a portfolio confident enough to receive them. The stress test that follows is therefore not an audit of strangers. It is a government examining, at scale, the systems it has chosen to run at scale. That examination is where trust is earned, and it begins on the next page.

REFERENCE
[1] Digital Dubai Authority, "State of AI Report 2026 entity submissions analysis," Dubai, UAE, 2026.



The Stress Test

The proof delivered, the pressure applied: five tensions, nine outside voices, and the choices that decide whether trust survives scale

Chapter 3 showed the systems in production. This chapter tests the thinking behind them: five tensions, argued by the companies and researchers who live with them.

A government that runs on artificial intelligence is choosing between competing goods every day: innovation against regulation, speed against safety, performance against explainability, centralization against autonomy, automation against human oversight. Chapter 3 showed what Dubai's architecture delivers in production. This chapter puts the thinking behind it under pressure, and it does so in the only credible way: by handing the pen to voices outside government and asking each to argue one tension from where they stand.

Their arguments, in their words

The rule of the chapter is independence. Each article is a contributed argument, written under a named byline, and the arguments are the contributors' own: the editorial team set the questions and the length, nothing more. Six pieces come from industry: etisalat, Google, OpenAI, Microsoft, Meta, and talabat, companies deploying these systems at global scale and in this city. Three come from the academy: the Mohammed Bin Rashid School of Government, the University of Birmingham, and RIT Dubai, institutions whose business is scrutiny. Agreement was not a condition of inclusion.

How to read this chapter

The six industry insights run first, each carrying its tension as a running tag: etisalat and Google open on innovation versus regulation, OpenAI takes speed and safety, Microsoft performance and explainability, Meta centralization and autonomy, and talabat closes the strand on automation and human oversight. The three academic perspectives then take the same ground back: the Mohammed Bin Rashid School of Government re-examines the tension the chapter opened with, Birmingham answers Microsoft from the research side, and RIT Dubai closes the chapter across a full spread with a framework for governing AI agents through oversight by design. Read in sequence, the nine pieces are a stress test of the doctrine this report builds toward: positions first, scrutiny second, and the hardest question held for last. The pressure begins overleaf, with etisalat.

1 INDUSTRY INSIGHT:

Innovation vs Regulation: Creating the Right Environment to Scale AI

Jason Kettle Digital and AI Transformation Director, etisalat

Innovation — Regulation

When governance is clear and built into the journey from the beginning, it can accelerate delivery.



From my perspective, innovation and regulation should not be positioned as opposing forces. The real question is how we create the right environment for organizations to innovate at speed, while maintaining the level of trust, security and accountability required to scale.

This is something we see first-hand at etisalat. We work with government and large organizations not only on defining their AI ambition, but increasingly on the harder part: taking AI from individual use cases and pilots into production and ultimately scaling it across the enterprise.

And this is where governance starts to matter.

When governance is clear and built into the journey from the beginning, it can accelerate delivery. Teams understand the boundaries they are working within, what needs to be protected, where decisions can be made quickly and where greater oversight is required.

A simple example is an internal AI assistant helping employees search policies, summarize documents or prepare meeting notes. The risk profile is relatively contained, and with the right controls around data access, privacy and model usage, there should be room to move quickly.

Compare that with an AI capability supporting traffic management, critical infrastructure, public safety, or a government service that directly affects citizens. The level of scrutiny must be different. In those environments, explainability, human oversight, cybersecurity, data sovereignty, and operational resilience become much more important.

The challenge comes when both use cases are treated in exactly the same way.

If every AI initiative follows the same approval path, regardless of its purpose, data, impact, or level of risk, we will inevitably slow innovation down. This is why I believe the next evolution needs to be towards more practical, risk-based governance.

We need clear guardrails, but within those guardrails organizations also need room to experiment, learn and move quickly. For lower-risk use cases, teams should be able to innovate with greater autonomy. As the level of risk, criticality, or citizen impact increases, so should the level of governance and human oversight.

We also see this in the way AI platforms are being built.

Ultimately, the objective should not be maximum innovation or maximum regulation. It should be responsible innovation at scale.

If every business unit selects its own models, builds its own controls and goes through its own security and governance process, scaling becomes difficult very quickly. A better approach is to establish trusted enterprise platforms, approved models, common data controls and reusable governance patterns. Once that foundation is in place, new use cases can move through the organization much faster because teams are not starting from zero every time.

The same applies to sovereign AI.

For certain government and regulated workloads, the question is not simply whether an AI model performs well. Organizations also need confidence around where data is hosted, who has access to it, how the model is operated and whether the underlying infrastructure meets the required level of sovereignty and security.

Getting those decisions right upfront can remove significant friction later in the program.

This becomes even more important as we move into the next phase of AI. We are no longer talking only about deploying individual models or copilots. We are moving towards AI agents, enterprise-wide AI platforms and increasingly sovereign AI environments where organizations have greater control over their data, models and infrastructure.

At that point, governance cannot sit around the technology as an additional layer. It needs to become part of the operating model.

That means bringing together business, technology, data, cybersecurity, legal and governance much earlier. It also means establishing common standards and platforms that allow organizations to reuse what has already been approved rather than restarting the governance journey for every new use case.

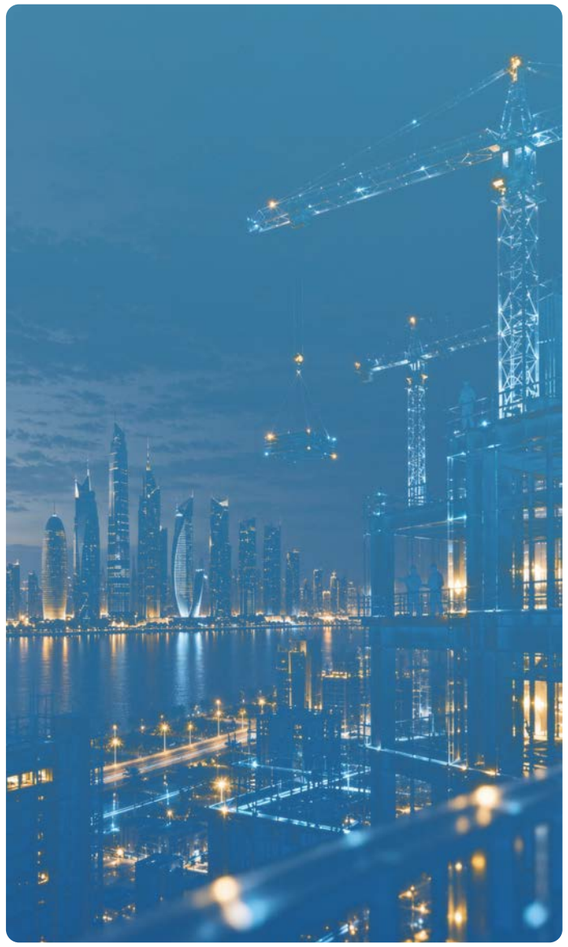
For me, this is where Dubai has a significant opportunity.

Dubai has consistently demonstrated an ability to move from ambition into execution. As AI adoption now moves towards city-wide scale, the opportunity is to create a governance environment that does not simply regulate AI but gives government entities and their partners the confidence to deploy it faster and more broadly.

Ultimately, the objective should not be maximum innovation or maximum regulation.

It should be responsible innovation at scale.

If we get that balance right, regulation becomes less about slowing technology down and more about creating the trust, clarity and confidence required to move faster.



If we get that balance right, regulation becomes less about slowing technology down and more about creating the trust, clarity and confidence required to move faster.

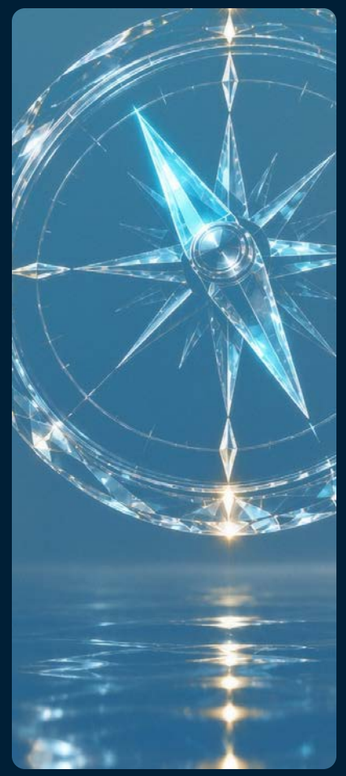
2 INDUSTRY INSIGHT:

Balancing Progress and Prudence: Innovation vs. Regulation in the Age of AI

Rasha El Hallak Head of Government Affairs and Public Policy, Google, UAE and the Levant

Innovation — Regulation

Our collective challenge is clear: designing governance frameworks that address genuine societal risks while unleashing AI's economic and developmental promise.



Artificial intelligence is the defining technological inflection point of our era. For the UAE, with its ambitious National Strategy for Artificial Intelligence 2031, the question isn't whether we govern AI, but how we do so without extinguishing its extraordinary potential. As policymakers, technologists, and industry leaders navigate this frontier, our collective challenge is clear: designing governance frameworks that address genuine societal risks while unleashing AI's economic and developmental promise.

Technological anxiety is not new. Virtually every foundational breakthrough (from the fifteenth-century printing press to mechanical textile looms, early electrification, and the commercial internet) was met with widespread apprehension regarding social disruption, workforce displacement, and moral order. Yet, time and again, history demonstrates that the greatest economic danger often lies not in technological progress itself, but in reactive, disproportionate over-regulation.

A notable historical precedent is the United Kingdom's Locomotive Acts, particularly the 1865 "Red Flag Act." Motivated by safety concerns and protective lobbying from horse-drawn transport operators, the law mandated that self-propelled road vehicles travel no faster than four miles per hour and be preceded by an individual carrying a red flag. While enacted under the banner of public safety, this rigid policy effectively paralyzed Britain's nascent automotive industry for over three decades, ceding technological, industrial, and economic leadership to Germany, France, and the United States. Preemptive, heavy-handed regulatory regimes consistently risk entrenching incumbents, suppressing domestic entrepreneurship, and triggering long-term productivity stagnation.

Rather than imposing blunt prohibitions, sustainable progress demands responsible innovation. At Google, our approach is defined by being both bold and responsible. We believe that technological capability and safety must advance in tandem. Responsible governance cannot be treated as a post-hoc compliance exercise; it must be built directly into the engineering lifecycle through proactive evaluation, red-teaming, open standards, and technical safeguards embedded by design.



Rather than imposing blunt prohibitions, sustainable progress demands responsible innovation.

A clear example of this approach in action is our development of SynthID, a pioneering watermarking technology. SynthID embeds imperceptible digital watermarks directly into AI-generated images, audio, video, and text. By establishing robust content provenance and detection mechanisms without degrading media quality or computational efficiency, such technical tools empower platforms, creators, and consumers to identify synthetic content and combat misinformation, fostering systemic trust while preserving the full utility of generative models.

This balanced paradigm is essential for unlocking the economic potential of AI, particularly for small and medium-sized enterprises (SMEs), which form the backbone of the UAE's non-oil private sector. AI fundamentally democratizes enterprise-grade capabilities that were once exclusive to large conglomerates. Local startups and SMEs are already leveraging AI tools to optimize complex supply chain logistics, streamline multilingual customer support in Arabic and English, automate routine administrative workflows, and accelerate product design at a fraction of traditional costs. In vital national priority sectors such as healthcare, clean technology, and trade finance, AI allows smaller enterprises to scale rapidly and compete on the global stage.

AI fundamentally democratizes enterprise-grade capabilities that were once exclusive to large conglomerates.

The UAE is already a global benchmark for progressive governance. By prioritizing regulatory sandboxes, agile policy, and public-private collaboration over rigid mandates, the nation is leading the way. Combining technical accountability with adaptive oversight will ensure AI remains a powerful engine for diversification, jobs, and shared prosperity.



The UAE is already a global benchmark for progressive governance.

OpenAI

3 INDUSTRY INSIGHT:

Speed and Safety Are Not Opposites

Rod Solaimani MENA Policy and Partnerships Lead at OpenAI

Speed — Safety

Real-world deployment creates a powerful feedback loop that advances both model capabilities and security.



The most important lesson I have learned from deploying advanced AI is that speed and safety are not opposites. Done well, they reinforce each other. Real-world deployment creates a powerful feedback loop that advances both model capabilities and security.

If governments wait for perfect certainty before adopting AI, they will miss transformative benefits: better public services, more productive workers, stronger education systems, faster healthcare workflows, and more resilient infrastructure. But moving too quickly without safeguards risks losing the public trust that makes adoption possible.

At OpenAI, we call this iterative deployment. We responsibly deploy increasingly capable tools, learn from real-world use, improve them, and repeat. There is a limit to what any lab, government, or regulator can learn through theory alone. Careful deployment reveals where models create value, where users need more control, where risks emerge, and where governance frameworks need to adapt.

This approach only works if safety is built into every stage. Our Preparedness Framework evaluates frontier models for severe risks, including cybersecurity, biological misuse, persuasion, and autonomy, and determines whether safeguards are sufficient before release. We combine rigorous internal and external testing with assessments from independent experts. Our System Cards explain evaluations, red teaming, and mitigations behind major releases, while inviting public feedback. Our Model Spec defines intended model behavior, including how models should follow instructions, resolve conflicts, respect user freedom, and remain within safety boundaries.

This is especially important for governments moving at Dubai's pace. Fast adoption requires partners who can explain not only what a model can do, but how it has been evaluated, where its limits lie, what safeguards are in place, and how those safeguards will evolve.

Red teaming is central to that effort. We work with internal and external experts before and after deployment, particularly in high-impact areas. We also collaborate with public institutions, including the U.S. Center for AI Standards and Innovation and the UK AI Security Institute, recognizing that governments bring expertise and accountability that companies cannot provide alone.

There are practical examples of this balance. In cybersecurity, frontier capabilities should reach defenders quickly while reducing the risk of misuse. Our Daybreak program gives approved defenders responsible access to more capable tools, while people remain in control of changes. We used Astra to examine our own systems and moved 25 per cent of our engineers from production work to fixing the vulnerabilities it found. We then built tools and playbooks to help customers strengthen their own defences.

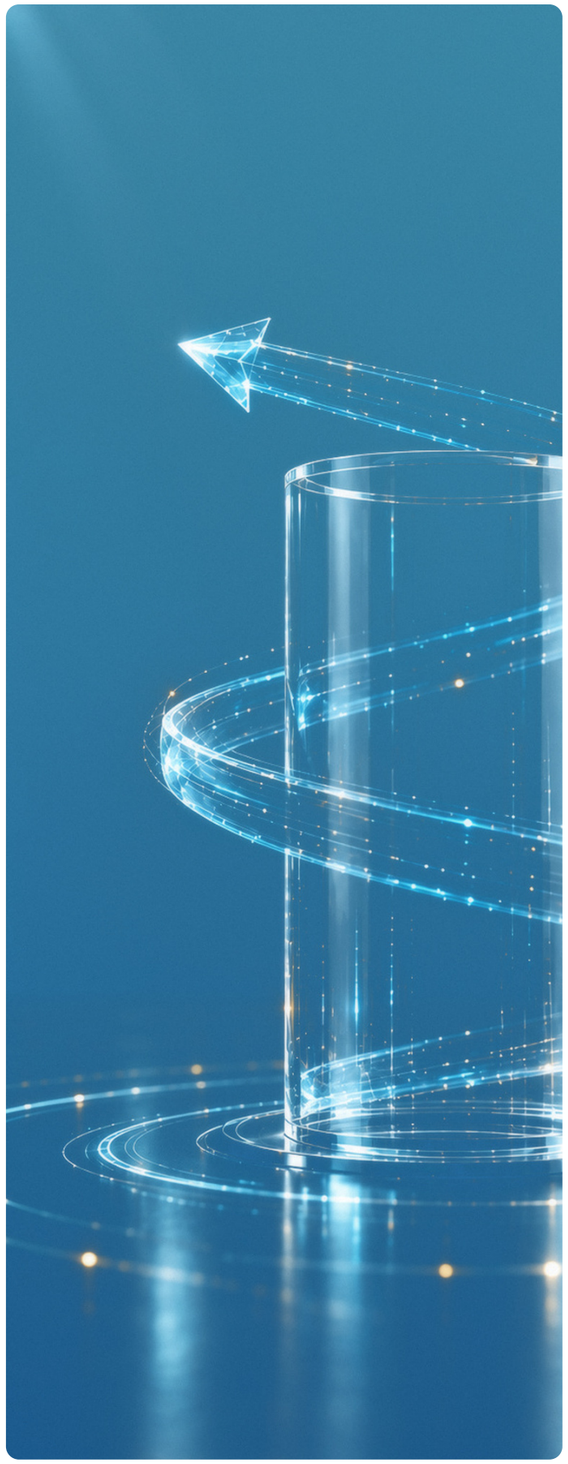
Youth safety demands the same thoughtful intention. Young people today are the first generation growing up with AI. They deserve experiences designed around their developmental needs. That principle guides ChatGPT for Teens. Users known or estimated to be under 18 receive stronger protections by default, alongside tools supporting learning and healthy use. Break reminders encourage time away, parental controls help families set boundaries, and, in limited high-risk situations, notifications can help connect a young person with offline support.

There is a limit to what any lab, government, or regulator can learn through theory alone.

Dubai has an opportunity to demonstrate what responsible scale looks like. The UAE has already set an ambitious vision for AI across government, energy, healthcare, education, transportation, and other critical sectors. The next phase will require urgency matched with discipline. Governments should expect AI partners to move quickly, but never casually: with evidence, transparency, red teaming, clear behavioral standards, strong safeguards, and a commitment to continuous improvement.



Youth safety demands the same thoughtful intention.



Dubai has an opportunity to demonstrate what responsible scale looks like.

4 INDUSTRY INSIGHT:

Performance vs Explainability: A False Choice for Governments Scaling AI

Yvonne Chebib Chief Commercial Officer, Health & Public Sector, Microsoft UAE

Performance — Explainability

In our experience, the most successful governments do not choose between performance and explainability. They recognize that both are essential.



Dubai's AI journey offers an important lesson for governments around the world: performance and explainability are not competing priorities.

Over the past decade, Dubai and the UAE have pursued one of the world's most ambitious AI agendas. They have invested early in digital infrastructure, established governance frameworks, built national AI capabilities, and encouraged innovation across both the public and private sectors. As a result, the UAE has emerged as a global leader in AI adoption, demonstrating how governments can move quickly while maintaining public trust. Microsoft's AI Diffusion research has consistently highlighted the UAE as one of the world's leading markets for AI adoption, reflecting both strong uptake and a favourable environment for innovation.

That achievement is particularly significant because it challenges a common assumption in the AI era: that organizations must choose between deploying the most capable technologies and ensuring those technologies remain transparent and accountable.

In our experience, the most successful governments do not choose between performance and explainability. They recognize that both are essential.

As AI systems become more powerful, governments naturally seek to maximize their benefits, whether by improving citizen services, accelerating decision-making, increasing productivity, or enabling entirely new forms of innovation. However, when AI begins influencing outcomes that affect people's lives, performance alone is not enough. Trust becomes just as important as capability.

Unlike many previous technology waves, AI has the potential to shape recommendations, automate actions and influence decisions at scale. When these systems support critical public services, stakeholders need confidence not only in the outcomes they produce, but also in the safeguards, governance mechanisms and accountability structures that sit behind them.



The next chapter of AI leadership will not be defined solely by technological capability.

Dubai's own experience illustrates why. The city's progress has been built not only on adopting new technologies, but on creating the policies, partnerships, governance frameworks and institutional confidence necessary to scale them responsibly. The result is a model where innovation and trust reinforce one another rather than compete.

As governments move from pilots and experimentation to large-scale deployment, explainability becomes less about understanding every technical detail of a model and more about understanding the systems around it. Can decisions be audited? Can outputs be traced to approved data sources? Are proper governance controls in place? Is there clear human accountability when decisions have significant impact?

Explainability is not simply a technical feature. It is an operational capability.

These questions become increasingly important as AI is embedded into everyday government operations.

The most successful AI implementations we see today share several characteristics. They provide traceability to source information. They maintain clear accountability between humans and machines. They incorporate robust testing and evaluation frameworks. They monitor performance and safety continuously after deployment. Most importantly, they keep people in control of high-impact decisions.

In other words, explainability is not simply a technical feature. It is an operational capability.

For government leaders, the goal should not be to make every AI model fully interpretable in every technical sense. Rather, the goal should be to build explainable systems: systems where outcomes are transparent, responsibilities are clear, risks are managed, and trust can be maintained over time.

This becomes even more important as governments begin deploying AI agents capable of taking action, not simply providing information. The more capable AI becomes, the more important governance, observability and accountability become. Success will increasingly depend on the ability to understand not only what an AI system recommends, but how it operates within established policy and operational frameworks.

Ultimately, explainability should not be viewed as a constraint on innovation. It is what makes innovation sustainable.

The theme of this year's report, Scaling with Trust, reflects a reality that Dubai has understood from the beginning of its AI journey: lasting progress comes not from innovation alone, but from creating the confidence to adopt innovation at scale.

As governments around the world seek to scale AI, Dubai's experience offers an important lesson: performance and explainability are most powerful when pursued together. High-performing systems create new possibilities, but explainability, accountability and governance create the trust that enables those possibilities to become enduring public value.

The next chapter of AI leadership will not be defined solely by technological capability. It will be defined by the ability to build systems that people trust, institutions can govern, and societies can confidently embrace. In that regard, Dubai is not only scaling AI, it is helping to shape the model for how AI can be scaled responsibly for the benefit of all.



Meta

5 INDUSTRY INSIGHT:

Choosing to Build Capability Rather than Rent It

Fares Akkad

Centralization — Autonomy

Centralization versus autonomy is not just a procurement decision, it is a question of where power sits once AI is embedded in the systems that run a city.



What Openness Actually Changes





Choosing open-source AI is choosing to build capability rather than rent it.

MUSE GLIMMER, THE OPEN-WEIGHT AGENTIC MODEL, AS THE ARTICLE DESCRIBES IT

- Apache 2.0 license
- Runs entirely on local hardware: a single consumer GPU or a laptop
- No cloud dependency
- Multi-step reasoning, tool use, complex task execution
- Every byte of data on the device

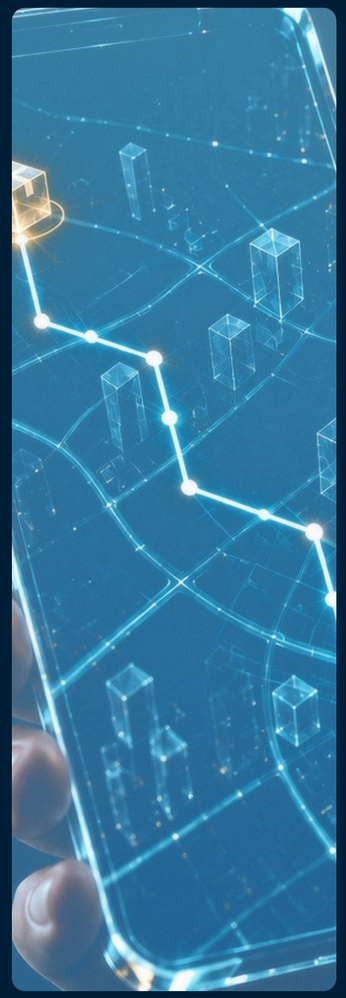
6 INDUSTRY INSIGHT:

Judgment sets the frame. AI finds the audience.

Sven Herzing Chief Technology Officer, talabat

Automation — Oversight

Humans set the frame. The model finds the audience inside it.



The question is usually framed as: how far can we automate before oversight becomes the bottleneck? In our experience, it runs the other way. Oversight is not the brake on automation: it is the reason we can let automation run at all.

What We Automate Fully, and Why

Personalization is fully automated. We make millions of ranking and recommendation decisions everyday and no human reviews a single one.The reason we can do that confidently is straightforward: the outcome is measurable, and any mistake is relatively small and quickly reversible.

What used to stay in human hands were the exceptions (a national holiday, a weather event, a match night) where past behaviour is a poor guide to the present moment. Reasoning models are closing that gap fast. Systems that can weigh context rather than match patterns are increasingly able to act on events they have never seen.

So the boundary of what we automate keeps moving. The interesting question is: what should not move with it?

When Automation Hit a Real Ceiling

We ran a series of experiments to automate the curation of our home screen, letting the system rank every component for every customer individually. It is the largest exposure surface we have, and the logic was sound: serve each customer exactly what they want right now. Individual relevance improved. Exposure across vendors and verticals declined.

The system was doing its job: optimizing the one thing it could measure, which was immediate customer intent. Nothing in it was defending the things it couldn't measure: which vendors get discovered, which categories get room to grow, what the platform should look like a year from now rather than this session.

That is not a modelling gap I can close with a better objective function. Deciding how much exposure a vendor deserves, balancing what a customer wants right now against what we deliberately choose to promote, is a judgment about whose interests a surface serves. Judgment cannot be automated, because it is a choice, not an optimum.

After several iterations, we landed somewhere specific: we use judgment and curation to define the experience, and AI to determine which customers that experience is right for. Humans set the frame. The model finds the audience inside it.

How a System Earns More Autonomy

Autonomy follows what is provable, not what is fast.

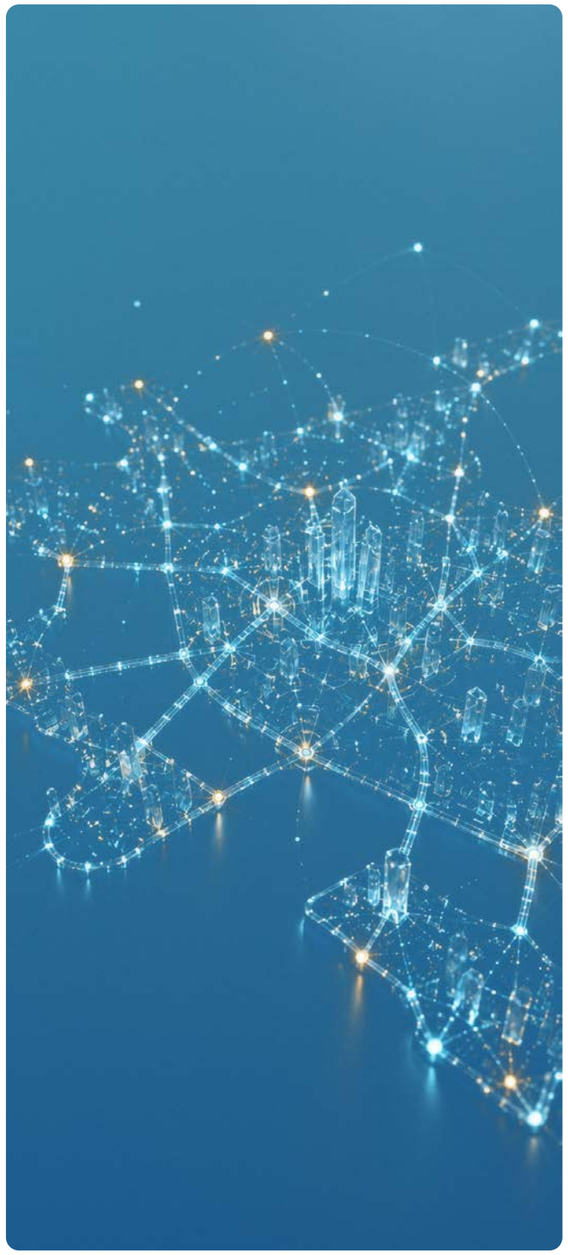
Autonomy follows what is provable, not what is fast. The practical test is two-part: is the outcome measurable, and can we afford small mistakes?

Relevance within a defined content slot passes both. We can measure it, and if the model gets it wrong for a customer today, the cost is low and the correction is fast, so we expand that autonomy and keep expanding it. Deciding how much exposure a vendor or vertical deserves fails both. The consequences compound slowly and are hard to reverse once the market has shifted. That stays human. The question we ask is not whether the model can make the call; it usually can. It is whether we would notice, quickly and clearly, if it made the wrong one.

There is a second reason to hold that line, and it is the one I think is underrated. These models are trained on the world's public knowledge. They are extremely good at the average, and by construction they pull toward it. Automate a curation process end to end and it converges on the safe middle. The output stops being wrong and starts being bland, which is the harder failure to notice, because no metric flags it.

Taste does not emerge from averaging. It comes from someone deciding what matters and having the conviction to steer the system toward it.

None of this is an argument for slowing down. We automate everything we can prove. We keep judgment where proof does not exist, and that is precisely what makes it safe to automate the rest.



None of this is an argument for slowing down.

THE AUTONOMY LADDER, IN THE ARTICLE'S OWN STRUCTURE

What We Automate Fully, and Why	When Automation Hit a Real Ceiling	How a System Earns More Autonomy
Personalization is fully automated.	We ran a series of experiments to automate the curation of our home screen, letting the system rank every component for every customer individually.	Autonomy follows what is provable, not what is fast.

7 ACADEMIC PERSPECTIVE:

Governing the Innovation-Regulation Tension in the Age of AI: Balancing Risks and Building Capability

Akmaral Orazaly and Fadi Salem Center for Future of Government, Mohammed Bin Rashid School of Government (MBRSG)

Innovation — Regulation

As AI moves rapidly across public services, with growing agentic and autonomous applications, the tension between innovation and regulation grows more prominent. Adoption without clear guardrails creates uncertainty, weak accountability and systems that are hard to scale. Control without implementation capability adds cost and delay without public value. The question for policymakers is: which forms of governance turn rapid adoption into lasting capability without undermining trust.



MBRSG's policy roadmap for AI governance in MENA finds that clear, inclusive and interoperable governance strengthens competitiveness by reducing uncertainty and enabling AI to scale responsibly [1]. Across hundreds of AI and digital companies in ten Arab countries, regulation appeared as both friction and enabler: a quarter of start-ups and SMEs linked it to higher operating costs, a fifth to slower innovation, yet a third said it made AI more secure and trustworthy and a similar share cited better market access [2]. In the UAE, nearly 70 per cent found AI regulatory uncertainty limiting and a quarter named clear, supportive regulation a leading government priority. These findings suggest that certainty in AI governance shapes planning, compliance and responsible AI deployment.

What matters, then, is not regulatory volume but regulatory design. Comprehensive risk-based systems provide protection and legal certainty, but can impose complex and costly compliance demands. Decentralized, sector-led approaches preserve flexibility, but can fragment requirements and accountability.

The GCC, for example, relies predominantly on soft regulation: strategies, ethical principles, sandboxes and institutional guidance alongside existing legal frameworks [3]. This preserves room for experimentation. However, the challenge is converting flexibility into durable capability through practical standards, interoperable rules, assessment mechanisms and audit capacity.

This implementation gap is not unique to the region: AI strategies are widespread worldwide, while evaluation, procurement support and delivery capabilities remain

uneven [4]. AI Governance instruments remain fragmented, and access to AI is of limited value without investment in skills, workflows and institutions [5].

Translating learning into practice that outlasts individual pilots is where the innovation-versus-regulation frame needs rethinking. AI governance has the dual functions of protection and enablement: mitigating risk and safeguarding public value as well as building the institutional capabilities responsible adoption requires. Weak governance not only fails to prevent harm and erodes trust; it also constrains innovation [6]. Innovation and regulation are interdependent features of the same system, not a trade-off.

Yet, that conversion is difficult. Value from innovation depends on an organization's capacity to learn from experience [7], [8], but AI maturity is slow to build as auditable processes, institutional trust and embedded accountability accumulate through practice, not through procurement [9]. Within government, this extends beyond training: as AI changes jobs, institutions must redesign roles, decision rights and accountability rather than leave adaptation to employees.

How can Dubai's institutions build sustained learning capacity while AI reshapes decision-making, measurement and accountability amid global fragmentation? As a first layer, public agencies need to invest in building the culture of ethical AI governance. Practically, this can be seeded through rolling out standardized and certified capacity-building programs for AI ethics (e.g., the Certified AI Ethics Professional program run by MBRSG and IEEE) [10]. A recent multistakeholder policy council convened by MBRSG's Center for Future of Government offers directions [11]. Four priorities emerge for Dubai's AI growth that are outlined in the box below.

Together, these measures treat learning capacity in AI implementation as an institutional capability to be built, rather than a high-level aspiration.

For Dubai's move to city-wide scale-up, the implication is not slower innovation but practical AI governance that lets institutions safely and confidently experiment, evaluate, correct and scale. The regimes that build lasting capability will not regulate the most or the least; they will let innovation move quickly, without moving faster than institutions can learn and adapt.

FOUR POLICY PRIORITIES FOR AI GOVERNANCE IN DUBAI · FROM THE ARTICLE

01 Separate AI oversight from experimentation

Adopt dual-track architectures in which performance monitoring and protected experimentation sit in distinct functions, so learning is not shaped by reporting incentives. Establish cross-entity 'learning councils' to compare pilots, identify transferable lessons and coordinate scale-up.

02 Measure AI trust and learning, not only performance

Complement KPIs with indicators of validated learning. For AI-enabled services, monitor trust indicators covering explainability pathways, grievance mechanisms and perceived fairness, reviewed on a fixed timeline to avoid metric lock-in and gaming.

03 Build AI-enabled learning infrastructure institutionally

Use AI to synthesize evaluations, pilots and research into decision-ready guidance and to surface patterns in partial failures. Require auditable decision-support logs for AI-assisted policy and high-impact service decisions, with explicit human judgment where outputs are contested.

04 Make honest reporting on AI safety, ethics and fairness "safe", especially at peak hype

Distinguish protected honest error from negligence, publish what changed, and reward cross-institutional knowledge sharing in performance and promotion criteria. Embed learning-loop training in leadership development and civil service curricula.

8

ACADEMIC PERSPECTIVE:

Beyond the Performance-Explainability Trade-off

Professor Kashif Rajpoot Professor of Medical AI, Interim Vice Provost (Education), University of Birmingham Dubai

Performance — Explainability

Successful deployment depends as much on how people understand, supervise and collaborate with AI systems as on model capability itself.



Artificial intelligence is increasingly embedded within public services, supporting decisions that affect citizens, public resources and critical infrastructure. As AI moves from research laboratories into institutional settings, the question of what constitutes a successful AI system is also changing. Much of the conventional academic debate has centred on balancing predictive performance with explainability, particularly in high-stakes applications where it is essential to understand how a model reaches its decisions [1]. This debate has played a pivotal role in advancing trustworthy AI. However, for governments deploying AI at scale, it is no longer sufficient. Public institutions are not evaluating models in isolation. They are adopting AI-enabled systems that must operate reliably, transparently and accountably within complex organizational and regulatory environments.

This shift changes the object of evaluation. A model that achieves state-of-the-art performance on benchmark datasets may still be unsuitable for public service if its outputs cannot be interpreted, challenged or integrated into human decision-making processes. Consider AI-powered business licensing, AI-assisted transport management, and AI-enabled healthcare auditing, all of which are already being deployed across Dubai's public sector. In each case, success depends on considerably more than predictive accuracy. Public officials need to understand when to rely on AI recommendations, when to question them, and how those recommendations fit within established governance processes. Human-AI interaction research demonstrates that successful deployment depends as much on how people understand, supervise and collaborate with AI systems as on model capability itself [2]. The key question for governments therefore extends beyond "Can this model perform the task?" to "Can this AI-enabled system support accountable public decision-making in practice?".

The answer lies in broadening the evaluation framework itself. Evaluating AI-enabled public services requires more than just balancing performance and explainability. It requires assessing whether AI can be deployed responsibly within the institutional context in which it will operate. Responsible AI provides this broader perspective, shifting attention from the characteristics of individual models to the readiness of AI-enabled systems for deployment. Explainability remains important, particularly where decisions affect individuals' rights and access to public services, but it is one of several characteristics that enable responsible deployment alongside robustness, fairness, human oversight and effective governance [3].

The future of public-sector AI will be determined by the ability of public institutions to translate high-performing AI models into trustworthy public services.

This broader perspective also changes how AI should be evaluated throughout its lifecycle. Governments must evaluate AI across the entire AI lifecycle, from model development through procurement, implementation, monitoring and ongoing governance. As Wirtz et al. argue [4], effective AI governance in public administration extends beyond algorithmic performance to encompass organizational processes, accountability mechanisms and human oversight. Evaluating AI therefore requires assessing whether the entire socio-technical system is prepared to operate responsibly in real-world public service environments.

This shift in evaluation has important implications as governments continue to expand AI across healthcare, education, justice and other public services. Success will increasingly depend on moving beyond narrow technical metrics towards a more holistic understanding of deployment readiness. For rapidly advancing digital governments such as Dubai, this perspective offers an opportunity to embed responsible AI principles as a foundation for innovation rather than as a constraint upon it. Ultimately, the future of public-sector AI will be determined by the ability of public institutions to translate high-performing AI models into trustworthy public services.

REFERENCE
[1] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," Nature Machine Intelligence, vol. 1, no. 5, pp. 206-215, 2019. doi: 10.1038/s42256-019-0048-x.
[2] S. Amershi, D. S. Weld, M. Vorvoreanu, A. Fourney, B. Nushi, P. Collisson, J. Suh, S. T. Iqbal, P. N. Bennett, K. Inkpen, J. Teevan, R. Kikin-Gil, and E. Horvitz, "Guidelines for Human-AI Interaction," in Proc. 2019 CHI Conf. Human Factors in Computing Systems, Paper 3, pp. 1-13, 2019. doi: 10.1145/3290605.3300233.
[3] A. Jobin, M. Ienca, and E. Vayena, "The global landscape of AI ethics guidelines," Nature Machine Intelligence, vol. 1, no. 9, pp. 389-399, 2019. doi: 10.1038/s42256-019-0088-2.
[4] B. W. Wirtz, J. C. Weyerer, and B. J. Sturm, "The dark sides of artificial intelligence: An integrated AI governance framework for public administration," International Journal of Public Administration, vol. 43, no. 9, pp. 818-829, 2020. doi: 10.1080/01900692.2020.1749851.

RIT Rochester Institute of Technology of Dubai

9 ACADEMIC PERSPECTIVE:

Automation with Oversight: Governing AI Agents Through Oversight by Design

Dr. Yousef Al Assaf President, RIT Dubai

Automation — Oversight

The goal is not maximum oversight. It is calibrated oversight that protects people without making useful automation impractical.



Artificial intelligence has entered a new phase. Earlier systems largely supported human decisions. Increasingly, autonomous AI agents can reason across multiple steps, revise plans, use digital tools, collaborate with other agents, and act with limited intervention [1]. For governments and organizations, the central question is how much authority should be delegated to systems acting on our behalf, and what must remain under meaningful human control.

This is one of the defining trade-offs of the AI era. Greater automation can improve productivity, speed, consistency, and service quality. Yet more autonomy can reduce opportunities for human review and make failures harder to detect, explain, or reverse. If every action requires manual approval, however, organizations may recreate the bottlenecks automation was meant to remove. The challenge is to establish a responsible relationship between automation and oversight.

Existing frameworks emphasize human agency, accountability, transparency, monitoring, and risk management [2]-[4]. However, keeping a human "in the loop" becomes less practical when agents execute thousands of interconnected actions at machine speed [5]. Reviewing every transaction is not practical. Oversight must move upward, from approving individual actions to governing objectives, permissions, risk thresholds, escalation conditions, and acceptable behavior.

This calls for Oversight by Design. Before deployment, an organization should determine what an agent may decide and execute, what evidence it must preserve, when it must pause or escalate, who may intervene, and who remains accountable. Unlike a checkpoint added after a system is built, Oversight by Design makes autonomy boundaries and human authority technically enforceable throughout the agent's lifecycle. Just as Privacy by Design embedded privacy into system architecture, oversight should become a core engineering requirement [6].

The appropriate level of oversight should reflect consequence, reversibility, uncertainty, and legal significance. Low-risk tasks may operate autonomously within monitored limits. Decisions affecting rights, safety, healthcare, financial entitlements, or access to public services should trigger stronger validation, explanation, appeal, and human review. The goal is not maximum oversight. It is calibrated oversight that protects people without making useful automation impractical.

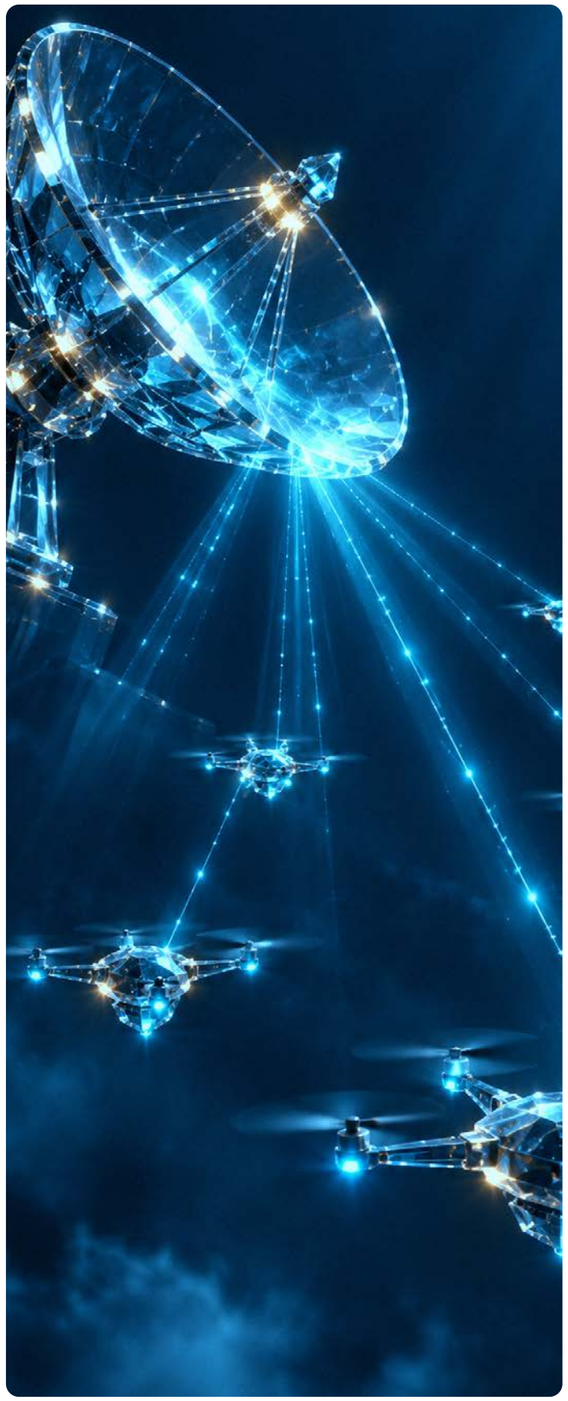
Consider a government building-permit process. An AI agent could verify documentation, assess compliance with planning rules, and approve routine applications. Its authority, however, should be explicitly bounded. Cases involving protected heritage, environmental impact, legal exemptions, conflicting regulations, safety concerns, incomplete evidence, or low confidence should automatically pause and escalate. The responsible official should receive the evidence considered, applicable rules, escalation reason, and a record of the agent's actions. Routine cases gain speed, while consequential cases retain meaningful human judgment.

Oversight is credible only when it can be demonstrated, tested, and improved.

For this model to work, every consequential agent needs a distinct and verifiable identity, a named organizational owner, and clearly bounded permissions [1]. An agent should not silently inherit the full access rights of the person or institution it represents. Its authority should be granted for a defined purpose, recorded separately, and capable of immediate revocation. These controls make responsibility visible and reduce the risk of an agent acting beyond its intended role without clear authorization or a way to stop it.

Oversight must also be measurable. Organizations should test more than accuracy and efficiency. They should examine whether agents escalate the right cases, whether humans respond in time, whether overrides work, whether decision records are complete, how quickly harmful actions can be contained, and whether essential operations continue when an agent fails [1], [2]. A policy promising human control is not enough if intervention cannot be exercised in practice. Oversight is credible only when it can be demonstrated, tested, and improved.

These safeguards carry real costs. Explainability, traceability, secure identities, escalation mechanisms, monitoring, and fallback procedures require engineering effort and ongoing investment. Some decisions will take longer. Yet weak oversight can create larger costs through erroneous decisions, cybersecurity incidents, regulatory penalties, litigation, public distrust, and the withdrawal of autonomous systems after visible failures. Oversight should be treated as enabling infrastructure for sustainable adoption, not as administrative overhead.



The future of AI will not be determined by how much society automates.

The human side is equally important. A person who formally retains authority but no longer understands the process cannot intervene effectively. Organizations must preserve expertise through training, simulated failures, manual fallback procedures, and clear responsibility for each consequential agent [1], [7]. Networks of agents may also produce unexpected interactions, diffuse accountability, and create risks that cannot be understood by examining individual models alone [8]. Good systems must know when not to act, and people must remain capable of taking command when those systems reach their limits.

Research universities have an essential role in examining these trade-offs independently. At Rochester Institute of Technology, work across human-aware AI, engineering, psychology, ethics, and human factors reflects a central principle: trustworthy AI depends not only on better algorithms, but on how people, institutions, and intelligent systems share authority [9]. Academic research can help governments and industry test where autonomy adds value, where oversight becomes essential, and how both can coexist without sacrificing innovation or public trust.

The future of AI will not be determined by how much society automates. It will be determined by how deliberately it decides what machines may do, what humans must retain, and how responsibility moves between them. The greater the autonomy of AI systems, the more sophisticated oversight must become. Automation may drive the next wave of productivity. Oversight by Design is what can allow that progress to scale safely, credibly, and with public confidence.



OVERSIGHT BY DESIGN · WHAT IS DECIDED BEFORE AN AGENT IS DEPLOYED

01	02	03	04	05
what an agent may decide and execute	what evidence it must preserve	when it must pause or escalate	who may intervene	who remains accountable

THE LEVEL OF OVERSIGHT REFLECTS

Consequence · Reversibility · Uncertainty · Legal significance

EVERY CONSEQUENTIAL AGENT NEEDS

A distinct and verifiable identity · A named organizational owner · Clearly bounded permissions

Framework as set out in the article; wording is the author's.

[1] Infocomm Media Development Authority and Personal Data Protection Commission Singapore, Model AI Governance Framework for Agentic AI, ver. 1.5, 2026.
[2] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, U.S. Department of Commerce, 2023.
[3] Organisation for Economic Co-operation and Development, Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449, amended 2024.
[4] European Union, "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)," Official Journal of the European Union, 2024.
[5] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation," IEEE Trans. Syst., Man, Cybern. A, Syst. Humans, vol. 30, no. 3, pp. 286-297, 2000, doi: 10.1109/3468.844354.
[6] A. Cavoukian, Privacy by Design: The 7 Foundational Principles. Toronto, ON, Canada: Information and Privacy Commissioner of Ontario, 2011.
[7] M. R. Endsley and E. O. Kiris, "The out-of-the-loop performance problem and level of control in automation," Human Factors, vol. 37, no. 2, pp. 381-394, 1995, doi: 10.1518/001872095779064555.
[8] S. Alqithami, "Adaptive accountability in networked multi-agent systems," in Proc. AAAI/ACM Conf. AI, Ethics, and Society, vol. 8, no. 1, 2025, doi: 10.1609/aies.v8i1.36536.

Chapter 04: The Stress Test
Academic Perspective: RIT Dubai





Dubai Emerging Doctrine for an AI-Enabled City

From a map of thirty-two postures to the decision framework Dubai now applies

Everything before this page states a claim or tests one. Chapters 1 to 3 set out what Dubai has built and how far adoption has travelled. Chapter 4 handed those claims to outside practitioners and scholars and asked them to press on five governance tensions: speed and safety, automation and oversight, innovation and regulation, performance and explainability, centralization and autonomy. This chapter draws the conclusions.

It opens with a map. Read as posture choices, the five tensions yield thirty-two theoretical combinations. Most cannot be held for long, and a handful cohere into recognizable ways of governing AI at city scale, each with its own trajectory, gains and costs. The map is a heuristic built for this report, and it is offered as one (5.1). The chapter continues with what a year of scaling taught the machinery that did it: five assumptions revised in public (5.2). It closes with the doctrine itself, Dubai's Emerging Doctrine for an AI-Enabled City: a decision framework for the city's government that states a governing objective, an operating principle for each tension, the evidence and controls leadership should expect before a system scales, the conditions under which the city will not proceed, and the review by which the doctrine will change (5.3). The editorial team signs off on the final page (5.4).

The word emerging is deliberate. The doctrine reflects Dubai's experience to date, and it will continue to evolve as AI capabilities, implementation evidence and public expectations change.

REFERENCE [1] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

1 Thirty-Two Ways to Scale

Five trade-offs, thirty-two combinations, a handful that cohere: a typology of how governments can hold AI, and what each posture costs



Positions are cheap. Combinations are destiny.

Every government scaling artificial intelligence answers the same five questions, whether it writes the answers down or not. How fast, against how safe. How much to automate, against how much to keep under human oversight. How far to let innovation run, against how firmly to regulate it. How hard to chase performance, against how much explainability to demand. How much to centralize, against how much autonomy to leave at the edge. Chapter 4 of this report put those five trade-offs, Speed vs Safety, Automation vs Oversight, Innovation vs Regulation, Performance vs Explainability, and Centralization vs Autonomy, in the hands of outside practitioners and scholars and asked them to press on Dubai's answers.

> *A posture is not five separate answers. It is one combination, and the combination, not the answers, sets the trajectory.*

WHAT THIS TYPOLOGY IS, AND IS NOT

The typology on these pages is a heuristic developed for this report. It is not an empirically validated classification, and no jurisdiction was scored against it. Its thirty-two cells are the logical combinations of two poles on five tensions. The three self-defeat patterns are the rule used to set aside combinations that contain their own contradiction. The five archetypes are the combinations that survive that rule and are internally consistent enough to be projected forward. Its value is as a decision aid: it makes the occupying of a posture a choice rather than an accident.

THREE PATTERNS OF SELF-DEFEAT

1. **The transparency trap**
   Full automation paired with an unconditional demand for explainability: the audit function cannot cash the cheques automation writes.
2. **Velocity without a vector**
   Hard speed paired with full edge autonomy and no central platform: motion mistaken for trajectory.
3. **The concentrated gamble**
   Centralized infrastructure paired with innovation-first permissiveness and no regulatory instrument: every entity inherits every other entity's riskiest decision.



This essay makes a different move. It treats each tension as a posture choice with two poles, which yields thirty-two theoretical combinations, and it asks a question that individual positions cannot answer: which combinations can actually be held? A government does not experience its tensions one at a time. It experiences them as a single operating posture, and the posture behaves as a system. A choice that looks prudent on one axis can quietly undermine a choice on another. The five answers interact, compound and occasionally devour each other, and it is the combination, not any single answer, that sets the trajectory: what the posture gains early, what it pays later, and what it becomes if nothing intervenes. This is why two governments with near-identical policies can end up in unrecognizably different places within five years, and why comparing individual policies across jurisdictions so often explains so little. The unit of analysis is wrong. Policies are ingredients; the posture is the recipe, and recipes fail as wholes.

Reading the map

The matrix on the following spread sets the five coherent archetypes against the five tensions, one row per posture, each cell showing the pole that posture takes. Behind those five rows stand all thirty-two theoretical combinations, and the map's first lesson is negative and it is the most important one: many of the thirty-two cannot be held for long. They are not wrong so much as unstable, combinations that contain their own contradiction and then scale it.

Three patterns of self-defeat account for most of the unstable cells, and together they are the rule by which the map was drawn. The first is the transparency trap: full automation paired with an unconditional demand for explainability. A government that automates decisions faster than it can explain them is writing cheques its audit function cannot cash; sooner or later it either quietly drops the explainability demand or quietly halts the automation, and the posture collapses into a different cell. The second is velocity without a vector: a hard speed commitment paired with full autonomy at the edge and no central platform. Everything moves and nothing converges; the government mistakes motion for trajectory and discovers, years in, that it has bought fifty sprints and no marathon. Speed is only meaningful as a property of the whole; a hundred fast parts can add up to a slow system. The third is the concentrated gamble: centralized infrastructure paired with innovation-first permissiveness and no regulatory instrument. Centralization concentrates whatever it carries; concentrate unexamined experimentation on shared rails and every entity inherits every other entity's riskiest decision. The cells that survive these patterns, the coherent postures, are few, and five of them are worth walking through. They are hypothetical constructs, deliberately so: no real jurisdiction matches any of them fully, and none is a portrait. Their value is that each is internally consistent, which means each can be projected forward without special pleading.

Five ways to scale that cohere: archetypes against the five tensions

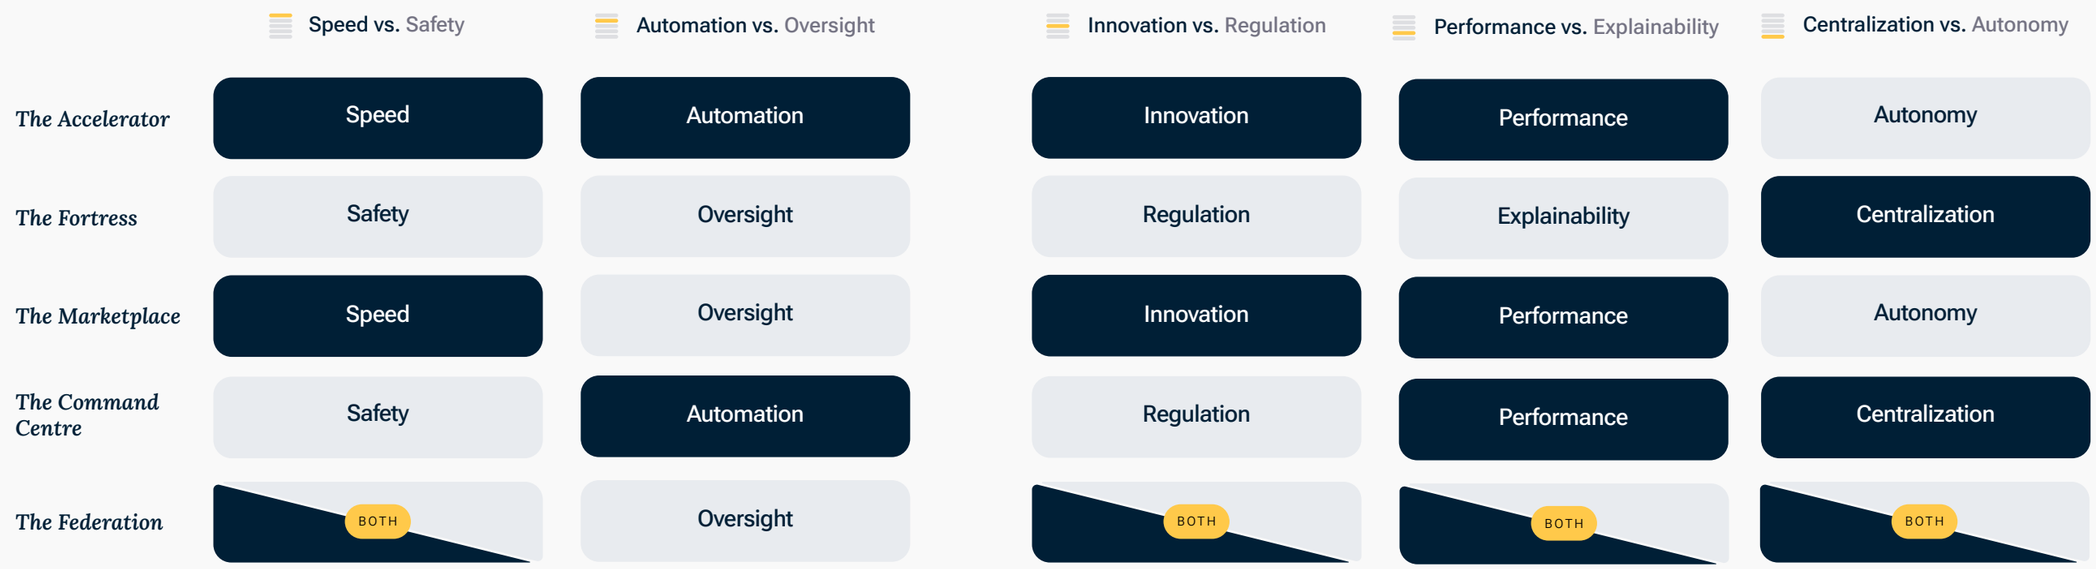


Figure 5.1: Five ways to scale that cohere: the archetype matrix. Five coherent postures out of thirty-two theoretical combinations; the rest fail one of the three self-defeat patterns described in the text. A heuristic developed for this report; all five archetypes are hypothetical constructs.

1 THE ACCELERATOR

Combination: speed, automation, innovation-first, performance-first, autonomy at the edge. The logic: the cost of being late exceeds every other cost, so entities are set loose and regulation is a promise for later. The gains are real and immediate: capability, momentum, a reputation for seriousness about the future. The losses arrive on a lag and compound: duplicated systems that cannot talk to each other, oversight that exists on paper, public trust that has never been tested. The trajectory bends at the first serious incident: with no assurance instrument in reserve, one failure can reprice a decade of wins. The Accelerator rarely dies; it panics into the Fortress, and panic writes worse rules than patience would.

> *"Automating at full speed with deferred oversight is not running a risk. It is running a countdown.*

2 THE FORTRESS

Combination: safety, oversight, regulation-first, explainability-first, centralized control. The logic: one catastrophe costs more than any velocity is worth, so every deployment is reviewed, every model explicable, every rule written before the tool it governs exists. The gains are legitimacy and auditability. The losses are quieter and larger: the queue becomes the signature, champions burn out, ambitious people and projects migrate, and so little is deployed that so little is learned. The trajectory is not collapse but irrelevance, perfect assurance over a shrinking share of what matters. Because nothing visibly fails, nothing forces the posture to revisit itself.

> *"The Fortress is the only archetype whose failure mode is indistinguishable, from inside, from success.*

3 THE MARKETPLACE

Combination: speed, human oversight held locally, innovation-first, performance-first, full autonomy. The logic: a government that trusts its entities and builds nothing central, so each entity runs its own race. The gains are real and local: experimentation flourishes, solutions fit their problems, and the best entities produce work of national quality. The losses are systemic and invisible from inside any one entity: infrastructure bought thirty times and compounded never, assurance so uneven that the resident's experience depends on which door they walk through, and a gap between the strongest and weakest that widens every year. The trajectory is an archipelago: islands of excellence with no mechanism that turns one island's learning into another's starting point.

> *"Its defenders point to the islands, rightly. Its critics count the ferries, and there are none.*

4 THE COMMAND CENTRE

Combination: safety, automation, regulation-first, performance-first, centralized everything. The logic: consolidation carried to its end, one platform, one rulebook, one owner, with automation at scale because central control makes scale cheap. The gains: compounding economics, consistency the resident can feel, and the only posture that can move the whole system at once, decisive in an emergency. The losses are structural: centralization concentrates failure as efficiently as capability, local knowledge is optimized out because a centre cannot see what it cannot aggregate, and one platform means one way of being wrong, replicated at machine speed. The trajectory is efficient fragility: years of quiet excellence, broken by rare, total failures that residents cannot understand.

5 THE FEDERATION

Combination: the first posture that refuses to answer every axis the same way everywhere. It centralizes the layer where compounding lives (platforms, infrastructure, assurance instruments) and federates the layer where judgment lives: what to build and how to use it. Oversight is constant, tiered to stakes; every other answer moves with risk. The gains: the Command Centre's compounding without its monoculture, the Marketplace's local fit without its archipelago. The losses: permanent coordination overhead, accountability seams between centre and edge, and a slower start: its instruments must exist before they pay. Held by instruments, not temperament, it can prove what it claims, and is the hardest posture to sustain. The trajectory: every pressure pushes it towards a purer neighbour, and a Federation that stops investing in its instruments drifts into one unannounced.

TRAJECTORIES ARE PREDICTIONS

Each portrait ends in a forecast, and forecasts are the only honest way to hold a typology: they can be wrong, and they say when they would be. These five trajectories are falsifiable claims about how postures behave over time, and this report intends to treat them that way. The map is an instrument, not an atlas of neighbours, and one property of it is not hypothetical: every government scaling AI occupies a cell, whether it chose the cell or drifted into one, and drift is a choice made with eyes closed. The purpose of a typology is not to grade the postures but to make occupying one a decision rather than an accident, with the costs priced before they are paid. The doctrine that closes this chapter states which posture Dubai has chosen, deliberately, and what the city accepts paying for it.

REFERENCE

[1] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

[2] Dubai State of AI Report 2026, Chapter 4 (this volume): the five tensions as examined by external contributors.



2 Lessons from the Scaling Year

Five assumptions the scaling year revised, and what changed because of them

> *A government that scales an exponential technology and reports no revised assumptions is not reporting candidly. These are ours.*

Scaling is a teacher with no syllabus. The architecture described in Chapter 2 and the adoption documented in Chapter 3 were not executed from a finished blueprint; they were built by an operating system that watched itself work and adjusted. This spread records five of those adjustments in a fixed form: what was assumed, what was observed, and what changed. None of them is a confession and none is trivial; an operating system that cannot name its revised assumptions is not learning, it is merely ageing. Together they are the strongest internal evidence behind the doctrine that follows, because a doctrine is only worth signing if it was written by an institution that has demonstrated it can change its mind.

LESSON ONE Readiness is uneven by design, not by failure

ASSUMED	OBSERVED	CHANGED
That government entities would climb the AI maturity curve at broadly similar rates once direction, policy and platforms were in place.	The spread was structural, not temporal. Entities with data-intensive mandates moved differently from entities whose value is judgment-intensive, and the city-wide aggregates (80 per cent of entities reporting an active AI governance board and 76.7 per cent mapping AI initiatives to strategy in the 2025 self-assessment) conceal honest variation in depth beneath them.	Enablement became tiered rather than uniform. The convening of Chief AI Officers across 27 entities, the AI+ program's 78-entity scope and CODE.AI's conversion pipeline each meet entities where they are instead of where a uniform plan assumed they would be.

LESSON TWO A trusted platform creates its own demand curve

ASSUMED	OBSERVED	CHANGED
That adoption of shared, sovereign AI infrastructure would build gradually and could be provisioned on a forecast.	Once entities trusted that, by the platform's design, prompts are not retained and models run under Dubai's control, demand compounded past every planning assumption: the running total of tokens processed on the shared inference platform grew by an average of 42 per cent a month across 31 government entities.	Capacity planning moved from forecast-led to architecture-led. The Infraboost program and the unified platform mandate of 1 April 2026 are, in part, the institutional answer to a demand curve that behaved exponentially while procurement behaved linearly.

LESSON THREE Self-assessment needs a counterweight

ASSUMED	OBSERVED	CHANGED
That entity self-assessment against the AI Integration Matrix Framework would be sufficient to steer city-wide priorities.	Self-assessment is indispensable and insufficient. It surfaces intent and effort reliably; it measures depth less reliably, and entities marking themselves hardest were often the furthest along.	Validation became a two-sided instrument. Central review now sits alongside self-assessment within the Matrix process, and the Dubai AI Seal adds an external tier of certification for the private-sector side of the ecosystem, so that no claim in the city's maturity picture rests on a single point of view.

LESSON Any instrument that asks an organization to describe itself must be paired with one that checks the description.

LESSON FOUR Governing applications is not governing agents

ASSUMED	OBSERVED	CHANGED
That the rulebook written for AI applications (policy, integration framework and certification) would extend naturally to agentic systems.	Agents broke the unit of governance. An application waits to be used; an agent acts, delegates and chains. The arrival of production agents in government back-office operations posed oversight questions, about accountability for actions rather than outputs, that application-level rules were never designed to answer.	In June 2026 Dubai completed the Federated Agentic AI Governance Framework and the Guidelines for AI Agents, extending the governance perimeter from what AI says to what AI does, ahead of the scale-up it expects rather than after it.

LESSON FIVE Capability is a workforce property, not a specialist property

ASSUMED	OBSERVED	CHANGED
That the scaling year's binding constraint would be technical, and that training a cadre of specialists would relieve it.	The constraint moved. Once platforms and policy were in place, the scarce resource was not the engineer who could build a system but the manager who could redesign a process around one, the director who could judge a risk tier, the non-technical employee confident enough to change how they work. Demand for AI capability appeared at every level of the hierarchy at once, and in Arabic as well as English.	Workforce development was rebuilt as a whole-of-government instrument rather than a specialist track. The AI+ program grew from a first cohort of 240 trainees across 36 entities into a five-track program spanning 78 entities, with a directors' track delivered in Arabic and a stated ambition of 50,000 trained by 2030. The revised assumption is now explicit: the city's absorptive capacity for AI, not its supply of AI, is the variable that decides the pace of everything else in this report.

REFERENCE

[1] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.

[2] Dubai Media Office, "Hamdan bin Mohammed directs the integration of all government services into a single unified AI-powered digital platform," mediaoffice.ae, 1 Apr. 2026.

[3] Digital Dubai Authority, "AI Integration Matrix Framework," whitepaper, Dubai, UAE, 2026.

[4] Digital Dubai Authority, "Federated Agentic AI Governance Framework" and "Guidelines for AI Agents," Dubai, UAE, June 2026 (cited in 2.11).

Five assumptions, five revisions, no casualties. The pattern beneath them is the same in each case: the instrument was not wrong, it was incomplete, and the correction added a counterweight rather than a reversal. That pattern, correcting by addition and at speed, is itself a finding, arguably the year's most important one, and the doctrine on the following pages is written above all to preserve it.

3 Dubai's Emerging Doctrine for an AI-Enabled City

Where Dubai sits on the map of thirty-two, the principle it applies to each tension, the evidence it requires before a system scales across the city, and the conditions under which it will not proceed



1 Purpose and status of the Doctrine

Dubai's Emerging Doctrine for an AI-Enabled City establishes how the city's government should make decisions as AI becomes more capable, embedded and autonomous in the services, systems and spaces of the city. It provides a common basis for deciding what to experiment with, what to deploy, what to scale and what not to pursue.

The Doctrine does not replace AI policies, regulations, technical standards or entity-level operating procedures. It provides the decision principles that connect them. Its purpose is to help leaders determine the appropriate position on five persistent governance tensions according to the public value, consequence and context of each use case.

It is written for the people who take those decisions: the leaders of government entities and the officials who prepare their decisions, the Chief AI Officers who carry the mandate inside each organization, and the coordinating institutions described in Chapter 2. Its subject is the government's own use of AI; its measure is what that use does for the city: the people who live in it, the businesses that work in it and the visitors who pass through it. It is applied through the instruments this report has described, the AI Policy for Government Entities, the AI Integration Matrix Framework, the Dubai AI Seal and the Federated Agentic AI Governance Framework: the Doctrine says how to decide, and the instruments say how to comply.

Four things should be clear from the outset. The Doctrine is a decision framework, not a technology manifesto. It does not prescribe the same answer for every AI system; it applies in proportion to consequence and autonomy. It is informed by evidence from implementation, including the lessons on the preceding pages and the challenge in Chapter 4. And it should be reviewed periodically, which is why it is issued as emerging and re-examined with every edition of this report.

HOW TO READ THE DOCTRINE

The chapter keeps three things apart. A doctrine principle is the enduring position. A decision rule is how the principle bears on a specific decision. An implementation instrument is the framework, standard, certification, registry or control used to apply it. Keeping them apart stops the Doctrine from becoming either too philosophical or too operational: it remains a concise leadership framework that connects clearly to the city's policies, standards and entity practices.

The map, signed: Dubai operates the Federation posture

	Speed / Safety	Automation / Oversight	Innovation / Regulation	Performance / Explainability	Centralization / Autonomy
The Accelerator	Speed	Automation	Innovation	Performance	Autonomy
The Fortress	Safety	Oversight	Regulation	Explainability	Centralization
The Marketplace	Speed	Oversight	Innovation	Performance	Autonomy
The Command Centre	Safety	Automation	Regulation	Performance	Centralization
The Federation	Both	Oversight	Both	Both	Both

Figure 5.3: The map, signed: Dubai operates the Federation posture.

Where Dubai stands on the map

The typology that opened this chapter ended with a promise: the map would be signed. Of the thirty-two postures, Dubai operates the Federation, and it does so by construction rather than by temperament. The platforms are central and sovereign; the judgment about what to build with them is federated to the entities that own the problems. Oversight is constant in principle and tiered in intensity. Innovation runs first, inside an envelope built from certification and framework rather than from prohibition. And on the two axes where a permanent answer would be a false one, speed and explainability, the posture is deliberately conditional: the answer shifts with the stakes of the use at hand. The instruments that hold the posture are named throughout this report: the unified platform mandate, the AI Policy for Government Entities of 2025, the AI Integration Matrix Framework, the Dubai AI Seal and the Federated Agentic AI Governance Framework completed in June 2026. The typology warned what this posture costs and how it fails: coordination is a permanent overhead, and a Federation that stops maintaining its instruments drifts into a purer cell without ever deciding to. A posture held by instruments can be audited; a posture held by temperament can only be asserted. The sections that follow are the audit standard.



2 Public value before technological novelty

One commitment stands above the five tensions: AI should earn its place in government through demonstrable public value to the city. The five tensions describe how Dubai manages the conflicts involved in delivering that value; the objective itself is not negotiable.

PEOPLE
better, more accessible and more complete services.

GOVERNMENT
higher productivity, institutional capability, resilience and decision quality.

ECONOMY
competitiveness, innovation, investment and responsible ecosystem development.

SOCIETY
inclusion, safety, trust and equitable access to benefits.

LEADERSHIP QUESTION

"What meaningful outcome will improve, for whom, and compared with what baseline?

Not every deployment must create a direct financial saving. Public value may take the form of improved service completion, faster or more consistent decisions, better accessibility, avoided risk, stronger government resilience or increased institutional capacity. Whatever its form, the value claim should be measurable and proportionate to the investment.

3 Five axes for deliberate decision-making

Dubai's AI posture is not defined by choosing one permanent position on each axis. It is defined by making those choices deliberately according to the value, consequence, autonomy and reversibility of each use case.

In some situations, the objectives at either end of an axis will genuinely conflict. In others, better technical design, institutional arrangements or assurance may improve both. The purpose of the Doctrine is to identify where a genuine conflict exists, determine the appropriate position and make the accepted costs explicit.

Each of the five axes that follow is set out in the same format: the tension, Dubai's operating principle, what it means in practice, the leadership question, the minimum evidence and controls, the accepted cost, and the conditions that would change the position. The format is the point. It makes the five positions comparable, and it makes them checkable.

4 AXIS 1: SPEED AND SAFETY

Move at mandate speed, with assurance proportionate to consequence

LEADERSHIP QUESTION

What degree of assurance is proportionate to the potential consequence and reversibility of this use case?

THE TENSION

Dubai seeks to adopt emerging AI capabilities quickly, but speed can increase exposure to operational failure, cybersecurity vulnerabilities, privacy breaches, poor-quality decisions, harm to customers and premature scaling. Safety does not always require slow delivery: common platforms, reusable controls and risk-tiered assurance can support both speed and safety.

WHAT IT MEANS IN PRACTICE

Government should:

- Move fastest where consequences are limitedand reversible.
- Increase testing and assurance as potential consequences rise.
- Establish control requirements before experimentation begins.
- Reuse approved platforms, components and assurance methods.
- Distinguish controlled experimentation from general deployment.
- Never weaken essential controls simply to meet a delivery deadline.

The one-year platform mandate is the principle in practice: an aggressive clock, running inside an architecture that was certified before the clock started.

MINIMUM EVIDENCE AND CONTROLS

Before scaling, leadership should expect:

- a defined outcome and baseline
- a risk classification
- reliability and security testing
- privacy and data-access assurance
- known failure conditions
- business-continuity arrangements
- rollback and incident-response mechanisms
- accountable authorization

ACCEPTED COST

High-consequence systems may take longer to deploy because stronger assurance is required. Dubai will sometimes be visibly slower than less demanding jurisdictions in exactly the domains where speed makes headlines, and it accepts that.

CONDITION THAT WOULD CHANGE THE POSITION

If assurance processes repeatedly cause delay without materially reducing risk, government should redesign the assurance process rather than remove essential safeguards.

5 AXIS 2: AUTOMATION AND OVERSIGHT

Automate execution, retain explicit accountability

LEADERSHIP QUESTION

What may the AI system decide or execute, what must remain human, and who is answerable when the system fails?

THE TENSION

Automation can improve speed, consistency and capacity. As AI moves from assisting and recommending to executing actions, responsibility can become unclear.

WHAT IT MEANS IN PRACTICE

Every significant AI system should define:

- What the system may recommend.
- What the system may execute.
- What remains subject to human authorization.
- What authority cannot be delegated.
- Who owns the system.
- Who is accountable for outcomes.
- When human intervention is required.
- How affected people can request correction or review.

Oversight should vary by system type. Human-in-the-loop: human approval is required before action. Human-on-the-loop: the system acts within defined boundaries under supervision. Human-over-the-loop: humans establish the mandate, monitor aggregate performance and intervene through predefined escalation mechanisms. No single model is appropriate for all applications. The Federated Agentic AI Governance Framework and the Guidelines for AI Agents (Chapter 2) carry this principle into agentic systems.

MINIMUM EVIDENCE AND CONTROLS

Minimum evidence and controls:

- a named government owner
- defined authority boundaries
- identity and permissions for agents
- decision and action logs
- escalation triggers
- human override and emergency pause
- correction and appeal mechanisms where people are affected
- periodic review of whether oversight is substantive or merely ceremonial

ACCEPTED COST

Meaningful human review may reduce the throughput of high-consequence systems. Dubai considers that a price, not a flaw.

CONDITION THAT WOULD CHANGE THE POSITION

Autonomy should not increase if government cannot identify failure, assign responsibility or intervene before material harm occurs.

6 AXIS 3: INNOVATION AND REGULATION

Enable experimentation through clear instruments; regulate deployment according to risk

LEADERSHIP QUESTION

What freedom is appropriate during experimentation, and what conditions must be met before exposure increases?

THE TENSION

Overly prescriptive regulation may constrain experimentation. Insufficient governance may expose people and government to systems that have not been adequately tested.

WHAT IT MEANS IN PRACTICE

Government should distinguish clearly between:

- Permission to experiment.
- Permission to deploy.
- Permission to scale.
- Permission to connect to shared government systems.
- Permission to make or execute consequential decisions.

An experiment may be acceptable even when city-wide deployment is not. The instruments available include risk classification, controlled environments and sandboxes, certification, technical and security testing, procurement conditions, impact assessments, deployment gates, licences or approvals where applicable, and post-deployment monitoring. They are examples of instruments, not proof that governance is effective merely because instruments exist. The Dubai AI Seal, the AI Integration Matrix Framework and the AI Policy's compliance regime are the instruments Dubai has in place; the escalation ladder of self-assessment, selective central review and stipulated timeframes is how they are applied.

MINIMUM EVIDENCE AND CONTROLS

Minimum evidence and controls:

- defined scope and duration of the experiment
- controlled user and data exposure
- approved operating boundaries
- pre-agreed deployment criteria
- monitoring of unintended outcomes
- independent review where consequences are significant
- an explicit decision to stop, redesign, continue testing or scale

ACCEPTED COST

Some experiments will not proceed to deployment, and not every investment in experimentation will produce a scaled system. Dubai accepts unscaled experiments as the price of unblocked ones.

CONDITION THAT WOULD CHANGE THE POSITION

If certification or other governance instruments do not predict real-world performance, prevent incidents or produce meaningful compliance, they should be revised rather than retained as procedural requirements.

7 AXIS 4: PERFORMANCE AND EXPLAINABILITY

Pursue the strongest performance that remains accountable for the intended use

LEADERSHIP QUESTION

What must the system explain, to whom, and at what level of detail for its use to remain accountable?

THE TENSION

The best-performing model may not always provide the level of explanation, traceability or contestability required for a particular government use. Equally, a highly explainable but weak-performing system may produce poor decisions and should not automatically be preferred.

WHAT IT MEANS IN PRACTICE

The required level of explainability should depend on:

- Whether the AI is advisory or determinative.
- The consequence of an incorrect result.
- Whether a person's rights or material interests are affected.
- Whether the decision can be challenged or reversed.
- Whether an official can understand and responsibly act on the output.
- Whether affected people can receive an adequate reason.

MINIMUM EVIDENCE AND CONTROLS

Minimum evidence:

- performance against a defined baseline
- a clear denominator and test population
- error types and their consequences
- performance across relevant languages and user groups
- known failure modes and limitations
- confidence or uncertainty information
- explanation appropriate to the government official
- explanation, notice or reason appropriate to the affected person
- a review and correction process

ACCEPTED COST

Dubai may decline a higher-performing system if it cannot provide the explanation, traceability or contestability required by the use case.

CONDITION THAT WOULD CHANGE THE POSITION

Improvements in explanation, validation or human-review methods may justify the use of more capable models in areas where deployment was previously inappropriate.

8 AXIS 5: CENTRALIZATION AND AUTONOMY

Centralize shared foundations; federate service judgment and accountability

LEADERSHIP QUESTION

Which capabilities benefit from government-wide consistency, and which decisions require entity-specific knowledge, ownership and judgment?

THE TENSION

Centralization can produce consistency, reuse, economies of scale, interoperability and stronger assurance. Excessive centralization can slow delivery, weaken entity ownership, reduce responsiveness to domain needs, create common bottlenecks and concentrate operational risk.

WHAT IT MEANS IN PRACTICE

In practice:

- Centralize where government-wide value compounds: shared infrastructure and compute; identity and access controls; interoperability standards; common interfaces; agent identities and registries; reusable assurance approaches; monitoring and incident mechanisms; common procurement requirements; government-wide measurement definitions.
- Federate where context and accountability matter: service design; use-case selection; domain-specific operating rules; engagement with customers and affected groups; operational judgment; entity-level change management; accountability for service outcomes.

The division of labour documented across this report, the centre building platforms and frameworks while twenty entities built the applications in Chapter 3, is this principle already in operation. It reinforces the Federation posture described in 5.1: neither centralization nor entity autonomy is treated as inherently superior.

MINIMUM EVIDENCE AND CONTROLS

Minimum evidence and controls:

- a clear division of responsibility
- common technical and data standards
- defined service-level expectations
- cross-entity escalation mechanisms
- an architecture exception process
- measures of platform adoption and reuse
- measures of duplication avoided
- evidence that shared platforms improve rather than delay delivery

ACCEPTED COST

Federation requires coordination and may create some duplication where local requirements legitimately differ. Convergence arrives more slowly than command would deliver it.

CONDITION THAT WOULD CHANGE THE POSITION

Where a shared platform consistently constrains necessary service innovation, or where entity-specific solutions create unacceptable fragmentation, responsibilities should be reassessed

Claim only what the evidence can support

The chapter keeps three things apart. A doctrine principle is the enduring position. A decision rule is how the principle bears on a specific decision. An implementation instrument is the framework, standard, certification, registry or control used to apply it. Keeping them apart stops the Doctrine from becoming either too philosophical or too operational: it remains a concise leadership framework that connects clearly to the city's policies, standards and entity practices.



9

From experimentation to scale

The five axes come together in one leadership decision gate. Before scaling a significant AI system, decision-makers should be able to answer the questions below. The gate is deliberately short: it is the minimum a leader should be able to answer, not the full assurance record behind the answers.

1 PUBLIC VALUE

- What need or outcome does the system address?
- What baseline will be used?
- Who benefits, and could any group be disadvantaged?

2 SPEED AND SAFETY

- Is assurance proportionate to consequence?
- Are the principal failure modes understood?
- Can the system be rolled back safely?

3 AUTOMATION AND OVERSIGHT

- What authority has been delegated?
- Who remains institutionally accountable?
- Can government intervene effectively?

4 INNOVATION AND REGULATION

- Has the system satisfied the conditions for deployment, not merely experimentation?
- Have legal, policy, procurement and technical requirements been met?

5 PERFORMANCE AND EXPLAINABILITY

- Does the system perform better than the relevant alternative?
- Can its operation and consequential outputs be adequately explained and challenged?

6 CENTRALIZATION AND AUTONOMY

- Is the capability appropriately shared or entity-specific?
- Will it support interoperability and end-to-end service delivery across the city?

7 ECONOMICS

- Does expected public value justify full lifecycle cost?
- Are projected savings distinguishable from realized financial benefit?

8 RESILIENCE

- Can government monitor, correct, override, transfer or stop the system?

THE GATE RESULTS IN ONE OF FIVE DECISIONS:

1. Continue controlled experimentation.
2. Deploy within a limited operating boundary.
3. Scale with specified conditions.
4. Redesign and reassess.
5. Do not proceed.



Five outcomes are more useful than a simple approve-or-reject decision, because most systems that reach the gate deserve neither a blank cheque nor a refusal; they deserve conditions.

10 Boundaries and non-negotiable conditions

Dubai Government will not deploy or scale an AI system when:

1. no meaningful public or operational outcome is defined;
2. the underlying data is materially unreliable for the intended decision;
3. accountability cannot be assigned;
4. consequential errors cannot be detected or corrected;
5. required human review is impractical at the intended scale;
6. affected people cannot access an appropriate correction or appeal pathway;
7. cybersecurity or privacy exposure remains unacceptable;
8. total lifecycle cost is not justified by expected value;
9. government cannot maintain continuity, intervene or exit;
10. the evidence does not support the strength of the claim being made.

A credible doctrine describes not only how Dubai enables AI, but also the circumstances in which the city will decline to use it.

> *"This document expires in one year, by design*

11 A living doctrine

The Doctrine should evolve through evidence rather than technological enthusiasm or institutional inertia. Changes in capability alone should not automatically change Dubai's position. A change should be supported by evidence of improved public value, performance, assurance, economics or accountability.

Each future edition of this report will identify which assumptions were confirmed, which were challenged, which principles changed, which controls proved effective or ineffective, which deployments were stopped or redesigned, and what new evidence changed Dubai's position on an axis. It will do so in a fixed form, the annual review table, so that the Doctrine's evolution is transparent and traceable.

Three tests are already on the record for 2027. If certification proves a poor predictor of real-world incidents, the third principle loses its foundation and prohibition-based regimes deserve a second hearing. If the coordination costs of federation visibly exceed the compounding gains of shared platforms, the fifth principle is wrong in its current form. And if tiered oversight fails to hold at agentic scale, if accountability seams open faster than they can be maintained once agents chain and delegate in volume, the second principle must be rewritten before volume arrives. The 2027 edition will open by answering all three with evidence, in the table on this page.

12 Closing commitment

Dubai's AI posture is not defined by choosing one pole of every axis. It is defined by making those choices deliberately and transparently according to the value, consequence and context of each use.

This Doctrine establishes how Dubai Government will move quickly where conditions permit, increase assurance where consequences demand it, preserve accountability as execution becomes more automated, enable innovation through clear instruments, require evidence before claiming impact, and combine shared government foundations with entity-level judgment.

Not every tension can be eliminated, and not every AI opportunity should be pursued. The objective is to ensure that increasing technological capability produces increasing public value for the city without weakening accountability, trust or institutional resilience.

Axis	Previous position	Evidence gathered	Observed outcome	Unresolved uncertainty	Revised position, if applicable	Resulting government action
Speed and safety	Mandate speed, assurance proportionate to consequence	*2027*	*2027*	*2027*	*2027*	*2027*
Automation and oversight	Automate execution, retain explicit accountability	*2027*	*2027*	*2027*	*2027*	*2027*
Innovation and regulation	Enable experimentation through clear instruments; regulate deployment by risk	*2027*	*2027*	*2027*	*2027*	*2027*
Performance and explainability	Strongest performance that remains accountable for the intended use	*2027*	*2027*	*2027*	*2027*	*2027*
Centralization and autonomy	Centralize shared foundations; federate service judgment and accountability	*2027*	*2027*	*2027*	*2027*	*2027*

REFERENCE
[1] Dubai State of AI Report 2026, 5.1, "Thirty-Two Ways to Scale" (this volume).
[2] Digital Dubai Authority, "AI Policy for Government Entities," Dubai, UAE, March 2025.
[3] Digital Dubai Authority, "AI Integration Matrix Framework," whitepaper, Dubai, UAE, 2026.
[4] Digital Dubai Authority, "Federated Agentic AI Governance Framework" and "Guidelines for AI Agents," Dubai, UAE, June 2026.
[5] Dubai Media Office, "Hamdan bin Mohammed directs the integration of all government services into a single unified AI-powered digital platform," mediaoffice.ae, 1 Apr. 2026.
[6] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025.



4 A Report Written to Be Revisited

A closing note from the editorial team



This report has tried to hold itself to the standard it describes. Scaling with trust means showing the working, and so these pages have done more than record a year of progress. Chapters 1 to 3 set out what Dubai built and what it measured. Chapter 4 invited practitioners and researchers to examine that record in their own words. Chapter 5 stated the posture the city has chosen and the decision framework behind it, in terms clear enough to be returned to.

Candour also means noting what remains open. Much of the evidence in Chapter 3 is reported by the entities themselves, and the report labels it as such rather than presenting it as independent audit. The measurement of explainability is a young science, and this report has said so rather than overstate it. The agentic era described here is barely a year old inside government, and the oversight model set out in these pages will mature alongside the systems it governs. A doctrine in its first edition is a beginning, not a settled conclusion. We believe the posture Dubai has chosen is the right one, and we have set out, in the preceding pages, the evidence by which each future edition will assess how it is serving the city.

That is the commitment this report makes to its readers: claims that can be checked, tensions that are stated plainly, and a doctrine reviewed every year at the pace of the technology it governs. The 2027 edition will open by returning to the questions this one has posed, and it will record what a further year of scaling has confirmed, refined or changed. We welcome the perspectives of readers, partners and peer cities in shaping that work, and we look forward to the conversation.

Until next year,

The Editorial Team, Dubai State of AI Report 2026

Sources and Hands

Every citation in one place, and the people behind the pages

Consolidated from the reference lists of all drafted articles, numbered in report order.
Article-level numbering is retained in each article; this list is the deduplicated master.

First cited in Chapter 1

[1] Stanford Institute for Human-Centered Artificial Intelligence, The 2026 AI Index Report. Stanford, CA, USA: Stanford Univ., 2026. [Online]. Available: https://hai.stanford.edu/ai-index/2026-ai-index-report (cited in 1.1, 1.4)

[2] Stanford Institute for Human-Centered Artificial Intelligence, "Economy," in The 2026 AI Index Report, 2026. [Online]. Available: https://hai.stanford.edu/ai-index/2026-ai-index-report/economy (cited in 1.1)

[3] Gartner, Inc., "Gartner predicts 40% of enterprise apps will feature task-specific AI agents by 2026, up from less than 5% in 2025," Press Release, Aug. 26, 2025. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025 (cited in 1.1, 1.4, 1.5)

[4] Gartner, Inc., 2026 CIO and Technology Executive Survey, 2025. [Online]. Available: https://www.gartner.com/en/articles/hype-cycle-for-agentic-ai (cited in 1.1)

[5] Middle East Institute, "AI, the Gulf, and the US: A primer," Washington, DC, USA, 2025. [Online]. Available: https://mei.edu/report/ai-the-gulf-and-the-us-a-primer/ (cited in 1.2)

[6] G42, "Global tech alliance launches Stargate UAE," Abu Dhabi, UAE, 2025. [Online]. Available: https://www.g42.ai/resources/news/global-tech-alliance-launches-stargate-uae (cited in 1.2)

[7] Middle East Institute, "From crude to compute: Building the GCC AI stack," Washington, DC, USA, 2025. [Online]. Available: https://mei.edu/report/from-crude-to-compute-building-the-gcc-ai-stack/ (cited in 1.2)

[8] Gulf International Forum, "Gulf AI infrastructure: Examining the business and economic case," 2025. [Online]. Available: https://gulfif.org/gulf-ai-infrastructure-examining-the-business-and-economic-case/ (cited in 1.2)

[9] UAE Cabinet, "Mohammed bin Rashid reveals new UAE Government framework to deploy agentic AI across 50% of government sectors' operations within two years," Apr. 2026. [Online]. Available: https://uaecabinet.ae/en/news/under-directives-of-uae-president-and-in-world-first-mohammed-bin-rashid-reveals-new-uae-government-framework-to-deploy-agentic-ai-across-50-of-government-sectors-operations-within-two-years (cited in 1.3)

[10] Government of Dubai Protocol, "Hamdan bin Mohammed appoints 22 Chief AI Officers across government entities in Dubai," 2025. [Online]. Available: https://www.protocol.dubai.ae/en/media-listing/news-events/hamdan-bin-mohammed-appoints-22-chief-ai-officers-across-government-entities-in-dubai/ (cited in 1.3)

[11] Gartner, Inc., "Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027," Press Release, Jun. 25, 2025. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027 (cited in 1.3, 1.5)

First cited in Chapter 2

[12] Dubai Media Office, "Hamdan bin Mohammed directs the integration of all government services into a single unified AI-powered digital platform," mediaoffice.ae, 1 Apr. 2026. (cited in 2.0b, 2.3, 2.9)

[13] Dubai Future Foundation, "Dubai Centre for Artificial Intelligence," dubaifuture.ae, accessed Jul. 2026. (cited in 2.1, 2.2)

[14] Dubai Media Office, "Hamdan bin Mohammed launches Dubai Universal Blueprint for Artificial Intelligence," mediaoffice.ae, Apr. 2024. (cited in 2.1)

[15] Dubai Media Office, "Dubai AI Week 2025 concludes," mediaoffice.ae, Apr. 2025. (cited in 2.1)

[16] Dubai Centre for Artificial Intelligence, "Dubai AI Seal," dub.ai/en/ai-seal, accessed Jul. 2026. (cited in 2.2)

[17] Zawya, "Dubai AI Seal certification coverage," zawya.com, 2025. (cited in 2.2)

[18] Arabian Business, "Dubai to unify all government services on one AI-powered platform," Apr. 2026. (cited in 2.3)

[19] Gulf News, "Dubai orders all government services onto a single AI platform within a year," Apr. 2026. (cited in 2.3)

[20] Emirates 24|7, "Dubai-it initiative announced alongside unified platform directive," Apr. 2026. (cited in 2.3)

[21] Digital Dubai Authority, "AI Policy," internal document (Policy Version 1.0), Dubai, UAE, 2025. (cited in 2.4)

[22] Digital Dubai Authority, "2025 AI Policy compliance self-assessment, aggregate results," internal document, Dubai, UAE, 2025. (cited in 2.4, 2.6)

[23] Digital Dubai Authority, "AI Integration Matrix Framework for government organizations," whitepaper, Version 1.0, Dubai, UAE, 2025. (cited in 2.5)

[24] Gulf Today, coverage of the public launch of the AI Integration Matrix Framework by Digital Dubai, 28 April 2026. (cited in 2.5, 2.6)

[25] Dubai Digital Authority, "AI Infrastructure Empowerment Platform: running total of tokens processed by month, September 2025 to October 2026," internal records, Dubai, UAE, 2026.

[26] Digital Dubai Authority, "CODE.AI Executive Council Resolution 44/2024 monitoring form," internal document, Dubai, UAE, 2025. (cited in 2.6)

[27] Digital Dubai, "AI+ Program launch release," Dubai, UAE, 15 April 2026. (cited in 2.6)

[28] Digital Dubai Authority, "AI Connect digest," internal document, Dubai, UAE, 2026. (cited in 2.6)

[29] Dubai Digital Authority, Shared Services Applications Department, "Smart Employee services: AI agents," internal document (Arabic), Dubai, UAE, 2026.

[30] Digital Dubai Authority, "CODE.AI: A network initiative to leverage AI and advanced data analytics for Dubai," internal document, Dubai, UAE, 2024. (cited in 2.7)

[31] Digital Dubai Authority, "Executive Council Resolution 44/2024 monitoring form: Centre for Data Intelligence and Advanced Analytics," internal document, Dubai, UAE, 2026. (cited in 2.7)

[32] Digital Dubai, "Digital Dubai launches 'AI+' to build an AI-ready government workforce," digitaldubai.ae, 15 Apr. 2026. (cited in 2.8)

[33] Digital Dubai, "Dubai Digital Skills Survey Results 2025," Dubai, UAE, 2025. (cited in 2.8)

[34] Digital Dubai Authority, "AI+ Program: HR orientation," internal document, Dubai, UAE, 2026. (cited in 2.8)

[35] Dubai Government Human Resources Department and Digital Dubai, "AI+ Program: Directors Track brochure (Arabic)," Dubai, UAE, 2026. (cited in 2.8)

[36] Digital Dubai Authority, "AI Infrastructure, Security, and Integration Standards," internal document, Dubai, UAE, 2026. (cited in 2.9)

[37] Digital Dubai Authority, "Infraboost: scalable AI infrastructure project deep dive," internal document, Dubai, UAE, 2025. (cited in 2.9)

[38] Dubai Media Office, coverage of the approval of the AI Infrastructure Empowerment Platform, mediaoffice.ae, Oct. 2025. (cited in 2.9)

[39] Gulf News, "Dubai Police deploy AI to make workforce planning smarter," gulfnews.com, 10 Sep. 2026.

[40] Dubai Digital Authority, "Financial Reconciliation service: launch announcement and user manual," internal documents, Dubai, UAE, 2026.

[41] Dubai Digital Authority, "SmartEmployee SmartPath: employee objective creation with AI assistant," internal presentation, Dubai, UAE, 2026.

[42] Dubai Digital Authority, "Agentic finance operations: AP invoice entry, AP payment entry and financial reconciliation," executive presentation, Dubai, UAE, Nov. 2026.

[43] Dubai Digital Authority, "AI use cases: AP Invoice Entry, AP Payment Entry and Financial Reconciliation agents," internal document, Dubai, UAE, 2026.

[44] Dubai Digital Authority, "GRP AI use cases 2026: ROI calculations," internal workbook, Dubai, UAE, 2026.

[45] Digital Dubai Authority, "Federated Agentic AI Governance Framework," internal document (v1.0), Dubai, UAE, 2026. (cited in 2.11)

[46] Digital Dubai Authority, "Guidelines for AI Agents," internal document (v1.0), Dubai, UAE, 2026. (cited in 2.11)

[47] Digital Dubai, "Dubai's State of AI Report," Dubai, UAE, 2025. (cited in 2.11, 2.13, 2.14, 2.15, 5.0b, 5.1, 5.2, 5.3)

[48] Digital Dubai Authority, "The AI Connect: intelligent collaboration from within," internal document, Dubai, UAE, 2026. (cited in 2.12)

[49] Dubai Government entity submissions, "State of AI Report 2026 entity data returns (private-sector partnership fields)," Dubai, UAE, 2026. (cited in 2.13)

[50] Dubai Media Office, announcement of the unified AI-powered government services platform directive, Dubai, UAE, 1 April 2026. (cited in 2.14)

[51] Digital Dubai, Dubai Government pavilion at GITEX Global, Dubai, UAE, October 2025. (cited in 2.14)

[52] Dubai Government entity submissions, "State of AI Report 2026 entity data returns," Dubai, UAE, 2026. (cited in 2.15)

First cited in Chapter 3

[53] Department of Finance, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.1a, 3.1b)

[54] Dubai Government Human Resources Department, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.1c)

[55] General Directorate of Identity and Foreigners Affairs, Dubai, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.1d)

[56] Dubai Public Prosecution, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.1e)

[57] Judicial Council of Dubai, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.1f)

[58] Awqaf Dubai, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.1g)

[59] Mohammed Bin Rashid Housing Establishment, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.1h)

[60] Roads and Transport Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.2a, 3.2b)

[61] Dubai Municipality, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.2a, 3.2c)

[62] Dubai Civil Aviation Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.2a, 3.2d)

[63] Dubai Civil Aviation Authority, "Director General's message on artificial intelligence," State of AI Report 2026 entity submission (supplementary), Dubai, UAE, 2026. (cited in 3.2d)

[64] Dubai Electricity and Water Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.3a, 3.3b)

[65] Dubai Health Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.4a, 3.4b)

[66] Dubai Health, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.4a, 3.4c)

[67] Dubai Health, "CEO message: AI as a foundational capability in an integrated academic health system," State of AI Report 2026 entity submission (supplementary), Dubai, UAE, 2026. (cited in 3.4c)

[68] Dubai Police, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.5a, 3.5b)

[69] General Command of Dubai Civil Defence, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.5a, 3.5c)

[70] Dubai Electronic Security Center, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.5a)

[71] Dubai Electronic Security Center, "Artificial Intelligence (AI) Security Policy, Version 1," Dubai, UAE. (cited in 3.5d)

[72] World Economic Forum, "AI and Cyber: Empowering Defenders," Geneva, Switzerland, 2026 (DESC contributing partner). (cited in 3.5d)

[73] World Economic Forum, "Global Cybersecurity Outlook 2026," Geneva, Switzerland, 2026. (cited in 3.5d)

[74] Dubai Customs, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.6a, 3.6b)

[75] Dubai Sports Council, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.6a, 3.6c)

[76] Community Development Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.7a, 3.7b)

[77] Dubai Culture and Arts Authority, "State of AI Report 2026 entity submission," Dubai, UAE, 2026. (cited in 3.7a, 3.7c)

[78] Digital Dubai Authority, "State of AI Report 2026 entity submissions analysis," Dubai, UAE, 2026. (cited in 3.10, 3.11, 3.12, 3.8, 3.9)

First cited in Chapter 5

[79] Dubai State of AI Report 2026, Chapter 4, "The Stress Test" (this volume): the five tensions as tested by external contributors. (cited in 5.1)

[80] Dubai Media Office, "Hamdan bin Mohammed directs the integration of all government services into a single unified AI-powered digital platform," mediaoffice.ae, 1 Apr. 2026. (cited in 5.2, 5.3)

[81] Digital Dubai Authority, "AI Integration Matrix Framework," whitepaper, Dubai, UAE, 2026. (cited in 5.2, 5.3)

[82] Dubai State of AI Report 2026, 5.1, "Thirty-Two Ways to Scale" (this volume). (cited in 5.3)



Contributors

This report was produced by the Digital Dubai Authority. It draws on submissions from the 20 government entities featured in Chapter 3, whose teams provided the implementation records, statistics and leadership statements that anchor the report's evidence base.

Acknowledgments

The editorial team thanks the AI Champions network, the entity liaison teams who coordinated twenty parallel submissions, and the reviewers whose comments sharpened every chapter.

دبـــــي الرقمية
DIGITAL DUBAI

Dubai's State of AI Report
Scaling with *Trust*

2026