

Horváth Middle East | Public sector

Advancing Government with Agentic AI

The Gulf is handing government work to AI agents. What that means for decision rights, workforce design and how progress is measured.

Saad Hamam

Partner and Managing Director, Horváth Middle East
September 2026

Two routes to agentic government. One constraint.

Research cut off 10 September 2026. Evidence standard: primary UAE and Saudi government sources, the International Labour Organization, the Center for Data Innovation and Horváth. First-hand observations are drawn from Horváth Middle East client work, anonymized and written as patterns.

50%



of UAE government sectors, services and operations move to agentic AI within two years. An operating model target, not a headcount target.

80,000



federal employees in the capability program approved by the Cabinet on 18 May.

66 / 100



Saudi Arabia's score in the Public Sector AI Adoption Index 2026, first of ten countries, with the highest daily public sector AI use in the ten-country index.

24 pts



the gap in Saudi Arabia between enthusiasm for AI (79) and enablement (55). Appetite runs well ahead of provision.

The argument in five lines

01

Two countries, two routes.

The UAE has set a deadline. Saudi Arabia built the layer underneath first and leads the ten-country public sector adoption index.

02

The binding constraint is authority, not the model.

If nobody on the program has the standing to amend the delegation of authority manual, the program has a ceiling, and the ceiling is a pilot.

03

The HR question is what people do the week after go-live.

Job architecture, performance measures and national entry roles must change before conversion.

04

Adoption is no longer the problem. Depth is.

Saudi appetite runs 24 points ahead of enablement.

05

Count volume, not services, and name where freed hours go.

Released hours only become value at a named destination, against a recorded before number.

The number is easy to misread.

On 23 April 2026, under the directives of President Sheikh Mohamed bin Zayed Al Nahyan, Sheikh Mohammed bin Rashid Al Maktoum announced that 50 percent of UAE government sectors, services and operations will move to agentic AI models for autonomous execution and decision making within two years. The framework calls for policies, processes and procedures to be redesigned around those capabilities, and states the principle it rests on: "People come first."¹

The number is easy to misread. It is not a published target to remove half of the federal workforce. On 18 May the Cabinet approved the implementation framework assigning roles to every ministry and federal entity, the first package of services for conversion, and a capability program for 80,000 federal employees across five occupational categories, from leadership through to the trainers who will deliver it.³ Mohammad Al Gergawi, Minister of Cabinet Affairs and Chairman of the National Committee for the Agentic AI Project, put it directly at the May retreat: AI does not replace people, it liberates them, giving them greater space to create and innovate.⁴

50%

of government sectors, services and operations, within two years

80,000

federal employees in the capability program, across five occupational categories

23 April

2026. The date the two-year clock started

Exhibit 1. What the UAE 50 percent target actually covers

What the target covers

- ✓ 50 percent of government sectors, services and operations
- ✓ Redesign of policies, processes and procedures
- ✓ A two-year timetable from April 2026

What it does not

- ✗ No published target to remove half of federal jobs
- ✗ Not a jobs target: a capability program for 80,000 employees sits alongside it
- ✗ Not a technology purchase: policies, processes and procedures are redesigned first

Source: UAE Cabinet, 23 April 2026, cross checked against implementation releases^{3, 6}.



"People come first."

The principle the framework states it rests on, and the reason the number is about how work is performed rather than who performs it. The same framework also states the program aims at reducing operational costs, boosting productivity and providing faster services.

How quickly is the UAE moving from policy to live agents?

Eleven days after the federal directive, Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum launched a parallel two-year initiative to move Dubai's private sector to agentic AI through Dubai Chambers.² On 1 September Dubai Chambers opened specialised agentic AI training to more than 14,000 member companies.¹⁰ That puts the same two-year clock in front of the commercial economy.

On 20 May, more than 400 ministers and federal officials met at the Agentic AI Retreat in Abu Dhabi, where the first four federal agents were launched for procurement, tax auditing, customer happiness and technical support.⁴ ADNOC reported more than 115 AI agents across HR, finance, procurement and auditing.⁴ Those are company figures, and the comparison has a limit. A company can change most of its own processes by internal decision. A ministry cannot change a process that is written into a regulation, and should not be able to.

11 days

between the federal directive and Dubai's private sector initiative

400+

ministers and federal officials at the Agentic AI Retreat, 20 May

115+

AI agents reported by ADNOC. Company figures, not government data

The first four federal agents, launched 20 May



Exhibit 2. Five months from target to second phase



Sources: UAE Cabinet, Dubai Media Office, Dubai Chambers, Emirates News Agency and the Federal Authority for Government Human Resources, April to September 2026.

How is Saudi Arabia building AI into government?

Saudi Arabia is moving in the same broad direction from a different starting point, and the comparison should not be forced. The Kingdom took a different route. Rather than a conversion deadline, it built the layer underneath first, and it has been at it since 2019.

SDAIA unveiled its AI Adoption Framework at the Global AI Summit in September 2024, alongside AI offices in 23 government entities.¹⁴ In July 2025 it launched a National AI Index assessing entity readiness across three pillars, seven dimensions and 23 subcategories.¹⁵ In March 2026 the Council of Ministers designated 2026 the Year of Artificial Intelligence.¹⁶ In July 2026 SDAIA issued a national AI risk management framework.¹⁸ SAMAI, the One Million Saudis in AI program run by SDAIA with the education and human resources ministries, passed one million certified citizens, and SAMAI 2 now runs with 11 ministries for government employees specifically.¹⁷

The measured outcome is the part most regional commentary has missed. In the Public Sector AI Adoption Index 2026, produced by Public First for the Center for Data Innovation from a survey of 3,335 public servants across ten countries, Saudi Arabia placed first with 66 out of 100, ahead of India and Singapore at 58. Among the 324 Saudi respondents, roughly two thirds use AI at work every day, the highest daily rate in the index, and the country ranked first on all five dimensions.¹⁹ That measures everyday use of AI tools rather than agents acting on their own, which is the step now beginning.

THE LAYER UNDERNEATH

Built from the bottom up since 2019, before any conversion target

- 2019**
SDAIA established. The national data and AI authority.
- Sep 2024**
AI Adoption Framework. AI offices in 23 entities
- Jul 2025**
National AI Index.
Three pillars, seven dimensions
- Mar 2026**
Year of Artificial Intelligence
- Jul 2026**
National AI risk management framework
- SAMAI**
One million certified.
SAMAI 2, eleven ministries
- Digital Government Authority**
National Cybersecurity Authority

Exhibit 3. Saudi Arabia ranks first of ten countries in the Public Sector AI Adoption Index



Index score out of 100. Source: Center for Data Innovation, Public Sector AI Adoption Index 2026¹⁹. The remaining six countries surveyed, Brazil, France, the UK, the USA, Germany and Japan, scored between 42 and 49.

WHAT THE UAE CAN BORROW

A measured baseline. Saudi Arabia scores entity readiness across three pillars and seven dimensions, so progress is visible before a deadline forces it. A comparable published readiness measure for UAE federal entities was not identified in the sources reviewed.

WHAT SAUDI ARABIA CAN BORROW

A dated commitment. The UAE directive names the services to convert, assigns roles to every ministry and entity, and puts a clock on it. Saudi Arabia has built the deeper foundation; a conversion deadline of the same kind has not been announced.

Each route is strongest where the other is still building, and the two together make a program that can be both delivered and measured.

Which decisions should an AI agent be allowed to make?

An agent differs from an assistant because it can be given permission to take actions. That moves governance from the quality of an answer to the authority to act. Retrieving a policy document is one kind of risk. Changing a record, approving an entitlement, initiating a payment or issuing an official communication is another. An agent that can act needs an authority boundary: what it may access, what it may decide independently, which actions require approval, when a case must escalate, how each action is logged, and who remains accountable for the outcome.

I would not draw that line by risk score. I would draw it with two questions. Does the decision create, remove or refuse a right or an entitlement belonging to a person? And is it contestable? You cannot appeal to a process that nobody authored. Somebody has to be answerable in the appeal, and “the system determined it” is not an answer any tribunal, or any minister, should accept.

THE TWO QUESTIONS



Does the decision create, remove or refuse a right or an entitlement belonging to a person?



Is it contestable?

If a citizen can appeal it, a person has to have decided it.

In countries where public servants have clear guidance and backing for AI use, 91 percent say they feel confident using it.¹⁹

Exhibit 4. The human and agent decision rights matrix



Proposed rule: automate the approval where criteria are met. Route every refusal to a person. Approvals and refusals of the same decision can sit at different tiers.

Source: Horváth analytical framework for discussion. Not an official UAE or Saudi government policy model.

What stays human

Refusals and revocations; anything decided on hardship or exception grounds, because the exception exists precisely because the rule failed; anything that sets a precedent, such as a first-of-kind interpretation of a regulation; enforcement, penalties and sanctions; and decisions about the entity's own employees.

Inside that boundary sits the cheapest good decision available to most entities. Approving a case that plainly meets the criteria and refusing one carry completely different consequences, yet most entities set one autonomy level for the whole decision. Automate the approval where the criteria are met, and route every refusal to a person. In most services the overwhelming majority of cases are straightforward approvals, so that single rule removes most of the risk at almost no cost in efficiency. Very few entities have made the distinction.



What mandatory human approval on one hundred percent of agent output actually asks of a person

1 in 100

cases a reviewer signs is the one that actually needed a person

4,000 cases a year.
40 of them matter.



Entities get the boundary wrong in both directions. Too permissive looks like agents given write access to the system of record early, because the integration happened to be easy. Too restrictive is more common and more expensive. Mandatory human approval on one hundred percent of agent output means a person reviews four thousand identical straightforward cases in order to catch the forty that matter. Review quality tends to erode quickly, approval becomes reflex, and the entity now has a weaker control than it had before, plus the cost of the agent. Horváth has documented the same pattern in supply chain planning, where formal human-in-the-loop control becomes symbolic once planners face more recommendations than they can challenge.²⁰

A person approving four thousand identical cases is not a control. It is a signature. Sampling, exception triggers and post-hoc audit are controls.

A case meets the criteria, or it does not

Approve, where the criteria are plainly met
Automate it. Most cases in most services are straightforward approvals.



Refuse

Route every refusal to a person. A citizen can contest it, and somebody has to be answerable in the appeal.



WHERE THESE PROGRAMS ACTUALLY STALL

Fourteen months later the process is still being done by hand.

What follows is a pattern drawn from earlier AI and automation programs in the region, and it is the pattern the UAE Agentic AI Project will meet. A large entity runs a well-executed pilot on a high-volume back-office process. The model performs. Accuracy is better than the manual baseline. Everyone in the room agrees it works. Fourteen months later the process is still being done by hand.

It does not stall at approval. It stalls at the point where the output has to enter the official record. Nobody asked, at design stage, whether an action taken by a system carried the same authority as an action taken by a named employee. When the question finally reaches legal, in month five, the honest answer is no, and there is no mechanism to make it yes, because the authority to act is defined in a delegation of authority manual that the program sponsor has no standing to amend. So the agent produces recommendations and a human retypes them into the system of record. Cycle time improves by nothing. Total cost goes up, because the entity is now paying for both.

How the stall actually looks



Two causes, and they usually travel together

Decision rights were treated as a compliance sign off at the end instead of a design input at the beginning.

The person accountable for the program had authority over the platform but not over the process.

The model is not your constraint. Your delegation of authority manual is."

The standard neither country has published yet

Both countries are building the governance layer. SDAIA's risk framework provides a common methodology for identifying, assessing, treating and monitoring AI risk, and the UAE has consolidated AI, data and digital government under a single authority and named its operating principle, human leads, AI enables.^{7, 11} Neither had published a task classification standard as at 10 September 2026 setting out which autonomous decisions are permitted in which services. Until one exists, entities should write their own and treat it as part of process design.

Month 5

when the question reaches legal, and the honest answer is no

0

improvement in cycle time. The agent recommends, a person retypes

Twice

the entity now pays for the agent and for the manual process

Eight things HR has to change before conversion, not after it

01

Plan in activities, not positions.

Public sector HR plans in positions and grades, tied to the organization structure. That instrument cannot express the sentence "this role loses forty percent of its task content and gains a supervisory function." Build an activity baseline by role family before conversion.

02

Redesign job architecture in parallel with the service, not behind it.

In most GCC public entities the job description is a formal instrument tied to grade, pay and civil service classification. If it says "processes applications" and an agent now processes them, the formal job no longer matches the actual job, with consequences for grading, evaluation and what the entity can lawfully ask of the person. In my experience this takes longer than the technology work.

03

Create the four roles most job catalogues do not yet define, grade them and give them a career path.


If they are not created, entities will fill them by informal assignment and lose the people who turn out to be good at them.

04

Confront the inversion of the competency model.

Value shifts from execution speed and accuracy towards judgement under ambiguity, exception reasoning and supervision of machine output. The highest performers in most processing functions today are the fastest and most accurate processors, and they are exactly the people whose task content disappears first. Without a deliberate move, an entity will demotivate or lose its best operators while retaining the profile it needed least.

Exhibit 5. Four roles most job catalogues do not yet define



Create them in the catalogue, grade them, give them a career path. Otherwise they are filled by informal assignment.

05

Decide the redeployment destination before go-live, and publish it.

Absorbing backlog, extending service hours, underserved services, inspection and field roles, an internal talent marketplace. The absence of an answer does not stay empty. It fills with rumour, and rumour damages adoption more reliably than any technical failure. Where staff believe the program is a redundancy exercise in disguise, the pilot fails for reasons that never appear in any report.

06

Move the national entry point up.

Entry-level and processing roles have historically been a major absorption channel for national graduates, and those are precisely the roles being converted first. That puts the AI target and national employment commitments into structural conflict inside the same entity, because government has been the main absorber of national graduates. The answer is not to slow the automation. It is to move the national entry point upwards, into the supervisory, exception handling and process owner roles, and rebuild the graduate program against the new role set. Very few have started, and this is where the two-year timetable will bite hardest. That needs a training path into those roles, and it needs to start now.

07

Rewrite performance measures before go-live.

If an employee's KPI is cases closed per month and the agent now closes the volume, the performance system is punishing exactly the behaviour the program requires. People will hold cases back from the agent to protect their numbers, and they will be right to, because the entity is paying them to.

08

Put HR on the body that approves autonomy levels, and convert HR's own processes first.

Recruitment screening, service requests, leave, letters and payroll queries are strong candidates, and any decision affecting an employee's status belongs in the human decides tier. Staff will read how the system treats them as the signal for how the entity treats people. And the workforce consequence of an autonomy level, meaning which role loses what content, is an HR fact decided today by technology and legal with HR not in the room.

The HR strategy question is not how many people we will need.

It is what those people will be doing in the week after the agent goes live, and whether anyone in the building owns that answer.

What will determine whether agentic AI actually delivers value?

Technology availability is only part of the answer. In the 7th Annual Horváth CxO Priorities Study, based on interviews with more than 1,000 CxOs across 32 countries and 16 industries, AI and digital transformation is the top strategic priority for 2026, ahead of cost and profit improvement for the first time. The same study scores the challenges limiting the AI transition on a four-point scale. Data foundation ranks first in both manufacturing and service industries at 3.12, the only barrier rated a major challenge across both clusters. Workforce capability and skills comes next, at 3.11 in services, just above the 3.10 threshold for a major challenge, and 3.05 in manufacturing, just below it.²³ The sample has no public sector cut, so treat it as a benchmark for the pattern rather than a finding about government. We see the same pattern in our government work: an agent cannot execute a reliable process on fragmented data, and a pilot cannot become an operating model while process ownership and decision rights stay where they were.

The same study shows that AI maturity is lowest exactly where agentic AI is aimed. In performance management and managerial decision making, 28 percent of companies have not started and 1 percent describe AI as fully integrated.²³

1,000+

CxOs interviewed across 32 countries and 16 industries


3.12

data foundation, the only barrier rated a major challenge in both manufacturing and services

1%

of companies describe AI as fully integrated in performance management and managerial decision making. 28 percent have not started

Exhibit 6. What executives say blocks a successful AI transition



Score on a four-point scale, 1 not important to 4 very important. Major challenge threshold 3.10.

Source: 7th Annual Horváth CxO Priorities Study, 2026, slide 28 (N=819). The chart shows the 819 respondents in the service and manufacturing clusters.

For the capability constraint there is direct public sector evidence, and it comes from the country that leads the index.

In Saudi Arabia, 84 percent of public servants say their employer gave them training to use AI tools. Half of those trained say it was more about compliance than about practical skills. Saudi Arabia scores 79 on enthusiasm and 55 on enablement, a 24-point gap between how much people want to use AI and how well their organization is set up for them to do it.¹⁹


84%

of Saudi public servants received employer AI training

50%

say it focused more on compliance than on practical skills

Exhibit 7. In Saudi Arabia, appetite runs 24 points ahead of provision



24

point-gap between enthusiasm (79) and enablement (55)

Index dimension scores out of 100. Source: Center for Data Innovation, Public Sector AI Adoption Index 2026, Saudi Arabia country factsheet.¹⁹

That matches what we see in our work, and the likely cause is structural rather than anyone's failure. Training is procured and reported as completions, so the delivery model optimizes for completions: a generic module, a certificate, a large number in a report. Nobody in that chain is measured on whether the work changed afterwards, so the work does not change afterwards. A 24-point gap between enthusiasm and enablement is not a motivation problem. Motivation there is the highest in the ten-country index.

It is a provision problem, and provision is a management responsibility, not a training vendor's.

Why training that reports completions does not change the work

A generic module

A certificate

A large number in a report

Nobody measured on whether the work changed

Exhibit 8. What useful capability building looks like

The unit is a workflow, not a tool

People train on their own entity's process, with their own data and exception cases in front of them, not on a general introduction to prompting.

01

It is differentiated by role

The supervisor of agent output, the exception handler, the process and rules owner, the AI assurance lead and the leader who signs off the autonomy level need five different curricula.

02

It builds supervision deliberately

Supervision is a new skill. A reviewer who has read forty correct outputs will approve the forty-first without reading it. That has to be trained and tested, ideally by seeding deliberate errors into a training version of the review queue and seeing whether they are caught.

03

A Horváth public sector case shows why workflow level measurement matters. A specialist public authority facing an annual review of tens of thousands of project descriptions used tuned large language models to classify them, with accuracy above 97 percent and a review process cut from months to a few hours.²² It was a classification task rather than autonomous decision making, but it had a baseline, a measurable outcome and a defined unit of work, which is what most AI programs cannot say. Across all ten countries in the index, 74 percent of public servants now use AI, while only 18 percent believe their government is using it very effectively.¹⁹

97%+

classification accuracy in a Horváth public sector case, with a baseline, a measurable outcome and a defined unit of work

74%

of public servants across the ten countries now use AI

18%

believe their government is using it very effectively

Adoption is no longer the binding problem in this region. Depth is.

THE TARGET SETS THE BEHAVIOUR, AND THE DENOMINATOR SETS THE TARGET

Illustrative


90%

of services converted



10%

of the transaction volume



is a report, not a transformation.

At the launch, Sheikh Mohammed bin Rashid Al Maktoum said the performance of ministers, directors general and entities will be assessed on their ability to adopt this transformation and their speed in implementing the new standards.¹ When speed of adoption is the measure, the denominator decides what gets adopted.

The risk is not the deadline. It is the denominator. Ninety percent of services and ten percent of the volume is a report, not a transformation.

I would not want the target removed. Without a date, few programs in this region leave the pilot phase. You get a permanent sequence of proofs of concept, each of them successful, none of them in production. A date creates the political cover to change a process, and changing a process is the expensive, unpopular part that everyone would otherwise defer. What a date needs beside it is a rule that no service enters the program without a recorded before state. A date is an excellent instrument for forcing decisions and a poor one for producing quality. The baseline is what stops the speed from eating the quality.

Count this

Share of transaction volume converted. Services that carry the workload. A recorded before state for every service that enters the program.

Not this

Share of services in the catalogue. Agents deployed. Use cases identified. Hours of training delivered.


How to calculate released capacity

This is the calculation entities keep asking for, and neither national program had published one as at 10 September 2026. The discipline is to separate hours released from value realized, because the first is an accounting entry and the second is the only thing a minister can defend.

Exhibit 9. From released hours to realized value

HOURS

BAR HEIGHTS ARE ILLUSTRATIVE



Report step 5 as the outcome. Steps 1 to 4 are the working that makes it credible.

An FTE equivalent can be derived for planning by dividing net redeployable hours by the entity's own productive hours per FTE. It is a planning figure, not a headcount forecast. Source: Horváth analytical framework for discussion.

Illustrative worked example

01 10,000 cases at 45 minutes each is 7,500 baseline hours.

03 Subtract oversight 600, transition load 400 and a surviving shadow process 500.

05 At 1,600 productive hours per FTE, that is about 2.8 FTE for planning.

02 Post redesign residual time is 1,500 hours, so the gross release is 6,000 hours.

04 That leaves 4,500 net redeployable hours.

06 Realized value is only what shows up at the named destination.

Suppose a redesigned process releases 6,000 hours in a year. If the service absorbs more demand, the value shows in throughput and waiting time. If people move to complex cases, it shows in backlog and quality. If no destination is named, the hours disappear into existing work and no benefit is ever visible. Horváth's AI to EBIT approach applies the same discipline commercially, holding every initiative to a stated contribution to earnings.²⁴ Government has no EBIT line. It has service outcome, cost per case and released capacity, and the logic transfers without modification.

If a minister only had one question: show me the before number. Where there is no before number, no claim about the program can be verified, including the favourable ones.

What should UAE and Saudi government leaders do next?

Seven things, in sequence. None of them requires a platform decision first.

01

Write the decision inventory for one high-volume service, before the strategy and before the platform.

the service owner, with legal
First 90 days

Who decides today, under what authority, what the rule is, what happens when it does not fit, the consequence of error, and where it is recorded. If the entity cannot complete it, that is the finding. *In the UAE, start with a service already inside the first conversion package, so the inventory feeds the two-year commitment directly. In Saudi Arabia, start where daily use is high and enablement is low, which is where the 24-point gap sits.*

02

Record the before state for every service entering the program: time, cost, quality, employee effort and service outcome.

the program office
First 90 days

That is what separates technology adoption from operating improvement.

03

Set decision rights in the design, not at sign off.

governance, with the process owner
First 90 days

For every process, state what the agent can decide, what needs approval, what escalates and who owns the outcome. Split approve from refuse while you do it.

04

Change the workforce plan from positions to activities, and start the job architecture work in parallel with conversion.

HR, with the process owner
First 90 days

It takes longer than the build, and it is the part that decides whether staff cooperate.

05

Make capability role-specific, and test supervision rather than attendance.

HR and the capability lead
First year

The Saudi evidence shows what happens when the volume of training outruns its depth.¹⁹ *In the UAE, use the five occupational categories in the 80,000 employee program as the role split rather than one curriculum. In Saudi Arabia, build on SAMAI 2, which already runs with eleven ministries, and make the next tier role-specific.*

06

Treat data ownership as a governance decision.

the entity data owner
First 90 days

The data problem in government is usually not a data quality problem. Two entities hold two versions of the same citizen record and nobody has the authority to decide which one is authoritative. No platform purchase resolves that.

07

Redesign the target operating model around the new allocation of work.

the head of the entity
First year

Horváth's Target Operating Model spans performance management, organization and structure, governance, processes, resources and capabilities, and technology.²⁵ Agentic AI touches all six.

What the next two years decide

Two years is long enough to convert services, and short enough to convert only the easy ones.

The governments that have something to show at the end of it will be the ones that settled three questions early. Which decisions an agent may take. What the people who did that work do next. And what counts as progress.

None of the three is a technology question, and all three are decided before the first agent goes live.

Horváth Middle East works with government entities on exactly these decisions.

Source register

Sources consulted for this paper. Bracketed numbers in the text refer to this list; entries without an in text reference were used to verify the program record. Source policy: primary UAE and Saudi government sources, the International Labour Organization, the Center for Data Innovation and Horváth. The Center for Data Innovation is a Washington based technology policy think tank; its 2026 index was researched by Public First and sponsored by Google, disclosed here in full. No research or statistics from competing management consultancies are used.

1. UAE Cabinet. 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. 23 April 2026. 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. Dubai Media Office. Hamdan bin Mohammed launches initiative to transition Dubai private sector toward agentic AI within two years. 4 May 2026. <https://www.mediaoffice.ae/en/news/2026/may/04-05/hamdan-bin-mohammed-launches-dubai-private-sector-shift-to-agentic-ai-within-two-years>
3. Dubai Media Office. UAE Cabinet, chaired by Mohammed bin Rashid, approves federal framework for Agentic AI Project implementation; 80,000 employee training program; first service package. 18 May 2026. www.mediaoffice.ae/en/news/2026/may/18-05/mohammed-bin-rashid-chairs-uae-cabinet-meeting
4. Dubai Media Office. Mohammed bin Rashid attends national retreat showcasing projects, key features of next phase of UAE Government's Agentic AI project: first four federal AI agents; ADNOC deployment figures. 20 May 2026. www.mediaoffice.ae/en/news/2026/may/20-05/mohammed-bin-rashid-attends-national-retreat-showcasing-projects
5. Federal Authority for Government Human Resources and MBZUAI. Strategic knowledge partnership for the agentic AI capability program. 21 May 2026.
6. UAE Government. UAE Government holds agentic AI workshop with participation of 50 federal entities. 10 June 2026. www.zawya.com/en/press-release/government-news/uae-government-holds-agentic-ai-workshop-with-participation-of-50-federal-entities-vkctpx9o
7. UAE Cabinet. Mohammed bin Rashid approves establishing the Artificial Intelligence and Data Authority. 14 June 2026. uaecabinet.ae/en/news/mohammed-bin-rashid-approves-establishing-artificial-intelligence-and-data-authority
8. UAE Cabinet. Ministerial Council for Artificial Intelligence and Development reviews plans and projects on transformation to agentic AI in federal government operations. 1 July 2026. uaecabinet.ae/en/news/chaired-by-mansour-bin-zayed-ministerial-council-for-artificial-intelligence-and-development-reviews-plans-projects-on-transformation-to-adopting-agentic-ai-in-federal-government-operations
9. www.gulftoday.ae/news/2026/09/09/ministerial-council-for-artificial-intelligence-and-development-reviews-second-phase-of-agentic-ai-project
10. Dubai Media Office. Dubai Chambers launches agentic AI training for more than 14,000 member companies. 1 September 2026. www.zawya.com/en/press-release/government-news/dubai-chambers-launches-specialised-agentic-ai-training-tracks-for-more-than-14000-business-group-and-business-council-member-companies-473930
11. Emirates News Agency (WAM). Launch of the strategic government work track under UAE Government 4.0 and the "human leads, AI enables" principle. August 2026.
12. UAE National Committee for the Agentic AI Project. Unified guide for designing government services powered by agentic AI. 22 August 2026.
13. Federal Authority for Government Human Resources. First cohort of the train the trainer program for the Agentic AI Project. 3 September 2026.
14. Saudi Press Agency. SDAIA launches the AI Adoption Framework at the Global AI Summit, Riyadh (12 September 2024); AI offices active in 23 government entities (1 December 2024). www.spa.gov.sa/en/N2170768 and www.spa.gov.sa/en/N2217019
15. Saudi Press Agency. SDAIA launches the National AI Index to measure government readiness for adopting AI. 23 July 2025. spa.gov.sa/en/N2365549
16. Saudi Press Agency. 2026 designated the Year of Artificial Intelligence following Cabinet approval. March 2026. spa.gov.sa/en/N2546742
17. Saudi Press Agency. SAMAI passes one million trained and certified citizens (9 November 2025); SAMAI 2 launched with 11 ministries (30 January 2026). spa.gov.sa/en/N2439614
18. Saudi Press Agency. SDAIA introduces a national framework for managing AI risks. 14 July 2026. spa.gov.sa/en/N2633781
19. Center for Data Innovation. Public Sector AI Adoption Index 2026, and the Saudi Arabia country factsheet. Researched by Public First, sponsored by Google. February 2026. datainnovation.org/2026/02/public-sector-ai-adoption-index-2026/
20. Horváth. Human on the Hook: why better AI recommendations do not automatically improve supply chain performance. July 2026. horvath-partners.com/de/media-center/artikel/human-on-the-hook
21. International Labour Organization. Workers' exposure to AI: what indicators tell us and what they don't. Research brief, 17 April 2026. ilo.org/publications
22. Horváth. Government agency accelerates data processes with large language models. Public sector client case. horvath-partners.com/en/media-center/featured-articles/government-agency-accelerates-data-processes-with-llms
23. Horváth. 7th Annual Horváth CxO Priorities Study, 2026.
24. Horváth. AI to EBIT: anchoring and scaling of AI. horvath-partners.com/en/core-offering/ai-to-ebit-anchoring-scaling-of-ai
25. Horváth. Target Operating Model / Future proof business and operating model. horvath-partners.com/en/core-offering/future-proof-business-operating-model

Author



Saad Hamam

Partner and Managing Director, Horváth Middle East
Saad.Hamam@horvath-partners.com

About Horváth Middle East

From strategy through execution, Horváth Middle East management consultants work with government entities, financial institutions, energy companies and defense organizations across the region to achieve measurable business impact through industry expertise and data driven outcomes.

We operate from our regional headquarters in Riyadh, with offices in Abu Dhabi and Dubai. Horváth is a global, independent management consultancy and a top consultancy for transformation, performance management and digitalization.

Offices

Riyadh Block 14, Building A, 2195 Prince Muhammad Ibn Abdulaziz St

Abu Dhabi The World Trade Center Mall

Dubai Dubai Mall Fountain Views

horvath-partners.com

Note. First-hand observations in this report are drawn from Horváth Middle East client work and are anonymized. All statistics are cited to the source register.