


RIYADH · KINGDOM OF SAUDI ARABIA

We turn what a government entity knows about its own technology into **an asset that entity owns - so that understanding stops leaving, whether the contract ends or the people who hold it move on.**



The headings are written as complete sentences, so reading only the bold lines gives you the whole argument in about sixty seconds.

This is a five-minute read. Nothing needs to be read in order, and most readers will need four of the thirteen sections.


I	1	Every entity buys the same understanding of its own systems, over and over.	3
THE COST			
II	2	We build a permanent model of how an entity's systems work and why - and the entity owns it outright.	5
WHAT WE BUILD, AND WHAT IT UNLOCKS	3	Every authorized engineer gets, on demand, the colleague who has been here fifteen years.	6
	4	The documentation and the audit evidence arrive as by-products of the work, not as write-ups afterwards.	7
	5	The System Brain is built once, and every project after it starts from the record instead of from discovery.	8
	6	You own every part of the platform, and your independence from us is measured on a schedule rather than promised in a clause.	9
III	7	The most capable AI models reason over your estate, and nothing real ever leaves your network.	10
WHAT YOU OWN, AND WHAT IS LAWFUL	8	One act here needs authority, and the regulator has already published the route.	12
	9	The lock-in nobody is defending against yet is cognitive, and we forbid it structurally.	13
	10	Nothing here asks you to adopt a new idea. Each part discharges an obligation you already carry.	14
IV	11	Five kinds of work we turn down, and each refusal protects something you are paying for.	15
WHO WE ARE, AND HOW TO CHECK US	12	This has to win the technical, the financial and the institutional conversation at once, and three principals hold one each.	16
	13	If you are looking for a reason to disbelieve this, here is the list.	18
CLOSING	→	The first step: one working session. You choose the task.	19

Every entity buys the same understanding of its own systems, **over and over.**

When a system is built, somebody comes to understand it completely: why it was designed that way, what depends on what, which failure means what, which shortcut was taken and why. That understanding is the most valuable thing the project produces, and it is the one thing that is never delivered. It leaves with the contractor, or with the engineer who takes another job.

The next change re-buys it. A new supplier spends the first weeks of an engagement working out how the system behaves. That work is real and it is not overcharging - but the entity has already paid for it, at least once, often several times. It never appears as a line item, because it arrives inside a day rate.

ONE ESTATE, THREE ENGAGEMENTS



 **Understanding, accumulating.** Formed by whoever does the work, and held only by them.

 **Bought again.** The first weeks of an engagement, spent working out how the system behaves.

 **Handover.** Nothing crosses it. The next engagement starts from zero.

That is why this cost is invisible rather than small.

The size of the line.

SAR 31.9 bn

GOVERNMENT ICT SPENDING, 2025 - DGA

55% TO 80% OF IT SUSTAINS WHAT ALREADY EXISTS, RATHER THAN BUILDING NEW - GARTNER

ROUGHLY HALF OF THAT IS CHANGE AND IMPROVEMENT, RATHER THAN REPAIR

SAR 9 to 13 bn

A YEAR, ALTERING SYSTEMS THE GOVERNMENT ALREADY OWNS

ONE PERCENTAGE POINT OF IT

SAR 90 to 130 million

a year, across government

And altering one safely requires **comprehending it first.**

What share of that is comprehension re-acquired rather than work performed is published by nobody, and it is unpublished for the same reason it is unrecovered: it is never booked.

Recovering it is not a saving on the work. It is the removal of a cost the government **pays twice, for knowledge it has already bought once.**

We build a permanent model of how an entity's systems work and why - and the entity owns it outright.

Three properties matter more than a feature list.


1 We call it a System Brain. It holds the intent behind each system, the reasoning that produced it, what depends on what, and the record of every failure and every fix.

It holds understanding, not telemetry.

A *digital twin* mirrors infrastructure in order to forecast its condition. A **cognitive twin** mirrors what the organization has learned about that infrastructure. Holding knowledge rather than measurements is what separates it from monitoring; holding reasoning rather than documents is what separates it from a repository.

2 It grades its own knowledge.

For every part of the estate it states whether what it holds is verified against source, inferred from documentation of uncertain freshness, or an acknowledged gap. You see exactly what it knows well, what it knows tentatively, and what it does not know. A search over documents cannot tell you that, and a supplier will not.



3 It stays current by construction.

From handover onward, one rule holds everything together: **no change to the estate without a design ticket raised with the System Brain**, enforced at the merge and deploy gates the entity already operates. A change that appears without one is treated the way the organization treats a bypassed code review - as an incident, not as drift to be reconciled later.

YOU OWN ALL OF IT

The underlying store

The map of what connects to what

The secrets vault

The encryption keys

The audit archive


03

MOVEMENT II
WHAT WE BUILD

OUR HALLMARK
SERVICE

Every authorized engineer gets, on demand, the colleague who has been here fifteen years.

Self-service agentic operations is our hallmark service, and the reason to hold a System Brain rather than merely to have built one. An engineer opens a session, describes a problem in their own words, and the System Brain investigates the live systems, designs the change, and carries it out - strictly inside the authority that person already holds. The person at the keyboard drives, decides and closes.



TWO THINGS CHANGE
IN THE FIRST MONTH

One

The queue in front of the one person who understands how something works disappears.

Two


Onboarding a new engineer becomes a session rather than a year.

Nothing autonomous touches production, at any authority tier. The authority boundary is the escalation mechanism: work exceeding an operator's tier stops and waits for someone who holds it.

The documentation and the audit evidence arrive as **by-products of the work**, not as write-ups afterwards.

THE METHOD

We build and modernize using frontier AI models under human authority, through a method we call **Governed Intent Compilation** - with every transformation recorded.



ITEM FOR ITEM

The method's second output matters as much as the first. Design intent, decision records, rejected alternatives, runbooks and verification records are emitted by the pipeline rather than written up later. They feed the System Brain. They are also, item for item, what the Kingdom's assessment instruments ask entities to produce:

Tested in real operating conditions rather than in a sandbox

Validated against the business need the technology was supposed to serve

Documented as outcomes occur


Resting on evidence rather than assertion

Entities are graded publicly and repeatedly - on the Emerging Technologies Adoption Readiness Index, on Qiyas, on DXMI, on Nadia - and the same artifacts are admissible to all four.

An entity that produces evidence continuously produces it once and submits it four times. An entity that assembles evidence when a cycle opens does that work four times a year.


The System Brain is built once, and every project after it starts from the record instead of from discovery.

That is the difference between buying a capability and buying a platform.



AND WITHOUT IT

Today each of these is bought as a separate project, and each one begins by **paying somebody to work out how the systems behave.**




EACH PROJECT STARTS AGAIN

You own every part of the platform, and your independence from us is **measured on a schedule** rather than promised in a clause.

No organization with long-standing processes leaps onto a new platform, and we never ask one to. The platform arrives through a bounded, paid engagement, time-boxed and priced fixed or capped, in which every performance claim we have made is scored against the entity's own historical records before handover into operation - and **every deliverable is the entity's property as it is produced**, so the decision de-risks at each step rather than only at the end.

What transfers is all of it: the underlying store, the map of what connects to what, the secrets vault, the encryption keys, the audit archive, and a record of how the estate works complete enough for your own engineers to work from. We call the arrangement **custody**.

Which leaves exactly one question worth asking, and we have written the arrangement so that you can fail us on it: **can your own people run this without us?**



- **A demonstrated export in open formats.** Performed, not documented in a clause.
- **The measured gap.** Your engineers resolve a sample of real work unassisted - before handover, then again at six, twelve and twenty-four months. The gap between assisted and unassisted performance is a reported number. If it widens, we are failing at the thing we were bought for, and the correction is priced into the standing arrangement rather than sold back as a remedy.
- **An annual witnessed cold-start test.** Cold because nothing of ours is present. Your own staff run live work end to end on a clean environment, using exported material alone.
- ▶ **A right to buy out.** After a stated period, purchase the operating model outright, end the standing fee, and run it internally.

You are not told you could leave. You practice leaving, on a date, with the result in your evidence pack.

The most capable AI models reason over your estate, and **nothing real ever leaves your network.**

Claude, GPT, Grok and Gemini all run outside any entity's perimeter, which is what puts every security committee in front of the same three options.

Refuse them

and accept a weaker system

Use them


and accept exposure

Or wait

while the mandate clock runs

We take **none of the three** as posed.

What leaves the perimeter is a faithful blueprint of the estate - its structure and the relationships between its parts, never its code - with every real name and address permanently swapped for a fake one, through a single audited exit and no other. Because the mapping is one-to-one and consistent, the structure the model needs in order to reason survives intact. The substitution table never leaves your premises; its keys are held non-extractably in your own hardware security module, and **no human role reads the mappings.**



Inside. The estate, the table and the keys. Everything real.

Crossing. Structure, with every name and address swapped.

Stopped at the wall. Code and secrets, with no policy option.

The property to test us on: [the crossing contains no judgment](#) - not a model's and not a person's.

Identical inputs produce identical outputs and identical audit records. **No human, however senior, approves a crossing.** There is no fast path and no emergency override, because the senior person who can wave something through is the failure mode the design exists to remove. Bulk access is architecturally impossible rather than administratively forbidden.

Four absolutes, stated without qualification

CODE AND SECRETS

Source code never crosses, at any classification, with no policy option and no enablement path. Secrets are hard-blocked and never waived.

PEOPLE

No identifier belonging to a natural person enters the system at all. No personal data is present by construction rather than by argument.

THE MODEL

No training, no fine-tuning, no self-modification of any model.

WHERE IT RUNS

No component of this platform is ever placed with a commercial cloud service provider, in or outside the Kingdom. It runs only on infrastructure the entity holds custody of - its own hardware, its own private cloud, or government cloud.

Where a classification forecloses the crossing altogether, the work runs on open-weight models inside your perimeter, at a [capability floor we publish rather than conceal](#) - and that you consent to at contract rather than discover during an incident.

SAID BEFORE YOU ASK IT

This is not a mechanism for lowering a classification. What permits a crossing is the classification of the content and the transfer rule attaching to it; the substitution does not create that permission, it bounds what can be read once the permission exists. And what lowers a classification is a determination made by your own Business Data Executive, not a transformation performed by a supplier. **A provider who offers de-identification as a route around a transfer rule is offering you a future audit finding rather than a control.**

One act here needs authority, and the regulator has already published the route.

Information originating inside the entity, in derived and de-identified form, crosses transiently to a private processor, and the result returns.

Two published instruments govern it, both issued by the National Data Management Office.

Data Classification Policy

NATIONAL DATA MANAGEMENT OFFICE


Data Sharing Policy

NATIONAL DATA MANAGEMENT OFFICE

NOTHING
TO INVENT

There is nothing to negotiate and nothing to invent. The transaction already has a defined form - a defined role for a private company, a defined route through your own Data Management Office, an instrument on the regulator's own template, and a signatory fixed by the classification of the material rather than by the value of the contract.

STATUTORY TEN-BUSINESS-DAY ASSESSMENT WINDOW



WHO DECIDES

Five determinations are made along the way. Every one of them is made inside your entity, by your own named officers. None is made by us. We specify a method and propose terms; we approve nothing. We can accept a determination or decline the work.

There is also a statutory ten-business-day assessment window attached to the route. It is a schedulable dependency and belongs in a project plan rather than being discovered mid-mobilization.

INSIDE
THE RULES

The regulator itself names masking and anonymization as permitted pre-processing methods, provided content is not affected. Because our substitution is deterministic and one-to-one, content preservation is machine-verifiable rather than asserted.


Which is why we are working **inside the published rules rather than around them.**

The lock-in nobody is defending against yet is cognitive, and **we forbid it structurally.**

Your contracts defend against being tied to a technology. Nothing at all defends against your operational knowledge accumulating inside a supplier's fine-tuned model.

TIED TO A TECHNOLOGY


YOUR EXIT CLAUSE



THE CLAUSE REACHES IT

ACCUMULATED IN A FINE-TUNED MODEL

YOUR EXIT CLAUSE



NOT PORTABLE, NOT INSPECTABLE


Where that happens, the knowledge cannot leave without abandoning what the entity has learned about its own systems. No exit clause reaches it, because the asset is neither portable nor inspectable: it cannot be exported in readable form and it cannot be audited.

The model does not learn. The corpus learns, and the corpus is yours.

No supplier accretes your knowledge, including us.

A RELATED TRAP,
ALREADY SITTING IN
THE INSTRUMENTS

The Kingdom's **Cloud First** direction, carried into DGA's *Digital Government Policies* at clause **8.2.9.1**, obliges entities toward cloud. Clause **8.2.3.1** of the same instrument forbids being bound to specific technologies. The common resolution - proprietary managed services from a cloud provider - satisfies the first while creating precisely the dependency the second exists to prevent, and the bill arrives years later under the legacy technology clause, **11.2.6.10**.



8.2.3.1 and 11.2.6.10 are the same clause separated by time.

We satisfy both, because here the choice of cloud is a deployment decision taken at the end rather than a dependency taken at the start. Open standards and open-source components throughout, no proprietary managed services, and moving is a redeployment rather than a rebuild.

Nothing here asks you to adopt a new idea. Each part discharges an obligation you already carry.

*Every clause below is from DGA's **Digital Government Policies**, and the numbering is the instrument's own, so each reference can be looked up and held against us. The 8.2 family is *Whole of Government Platforms*, the 9.2 family is *Digital Services Lifecycle and Upskilling*, and the 11.2 family is the *Technology Policy*.*

11.2.2.4

Institutional knowledge stops leaving with the contract

Results of all lifecycle stages documented and retained as a reference for continuous improvement.

9.2.4.1 to 9.2.4.4

Key-person risk removed, and the improvement measured rather than asserted

Annual plans to build digital skills and procure technologies **and knowledge assets**. *Knowledge asset* is a named procurement category in the regulator's own vocabulary, which is what lets a finance function book the System Brain as an asset rather than expense it.

8.2.3.1 11.2.2.6 11.2.6.8 11.2.6.10

Freedom of movement between technologies and clouds

Freedom from being bound to specific technologies, open-source practice, interoperable open standards, and a comprehensive approach to legacy technology.

11.2.3.3

Reuse instead of rebuild

Re-use over single-use custom build.

11.2.4.4

Proof becomes a by-product

A comprehensive methodology to measure the impact of emerging technologies.

11.2.6.5 11.2.2.7

Operations become self-service across the estate

IT service management with support to all users; technologies managed as a reusable government resource.

Five kinds of work we turn down, and each refusal protects something you are paying for.

These are commitments rather than caveats. The reason for each sits beside it.

01

No autonomous production change, at any authority tier.

The person at the keyboard closes the ticket, because supervision that holds is what makes systems of this kind safe.

02

No staff augmentation, no resale, no reseller margin.

We sell outcomes against a signed baseline, never hours - because a firm paid for time has no incentive to remove the queue.

03

No advisory work over an estate not under custody.

Advice without the record is opinion.

04

No monitoring product, no ticket system of record, no general-purpose assistant.

Your existing tools remain the systems of record, and the platform works through them.

05

No frontier answer on content a classification will not let cross.

And no degraded local substitute presented as an equivalent one.

12

MOVEMENT IV
WHO WE ARE

SEVENTY-NINE YEARS,
COMBINED

This has to win the technical, the financial and the institutional conversation at once, and three principals hold one each.

Tutelex is a Saudi company in Riyadh. Three principals, seventy-nine years of combined senior experience, and the reason that matters is that the three conversations tend to arrive in the same meeting rather than in sequence.



Yousuf Al Qasir

THE INSTITUTIONAL CONVERSATION

Twenty-five years founding and structuring commercial enterprises, with more than twenty-five public sector investment contracts managed and executed, and direct working knowledge of the Kingdom's procurement and regulatory machinery.

25

YEARS



Yasser Al Masri

THE TECHNICAL CONVERSATION

Twenty-six years in enterprise technology and AI systems architecture, latterly lead architect on enterprise AI systems for national operators across energy, refining, water, healthcare and insurance - at document and transaction volumes in the millions, under in-Kingdom sovereignty constraints throughout. Previously practice leadership at the Kingdom's sovereign cloud operator and at two systems integrators, and integration architect at the Kingdom's second-largest telecom operator.

26

YEARS



Khaldoun Shatat

THE FINANCIAL CONVERSATION

Twenty-eight years of executive experience in corporate finance, strategic planning, governance and program management, on programs valued in the hundreds of millions of riyals.

28

YEARS

Most firms hold one of those three at senior level. This team holds all three.

And nothing above is asked on trust.

What we offer is measurement on your own material.

SCORED

Every performance claim we make is **scored against your own historical records** before anything is handed over into operation.

TRANSFERRED

Every deliverable is **your property as it is produced.**

SELF-TESTED

Our own engineering practice runs on the same System Brain under the same custody discipline, **so the annual cold-start test is something we do to ourselves before we ask you to do it to us.**

FREQUENTLY OVERLOOKED

Under DGA's own instrument, a provider is itself an addressee of the policies, and the entity remains responsible for full compliance by any provider it contracts.

A provider's compliance posture is therefore your exposure rather than a negotiating point

- which is why this profile settles classification, the lawful route, the security architecture and the boundaries above in advance rather than during mobilization.

If you are looking for a reason to disbelieve this, **here is the list.**

RELIABILITY

Agents are unreliable, and demonstrations prove nothing

Agreed. Reliability rather than capability is the binding constraint, and a system that succeeds once will often fail on a re-run of the identical task. We report distributions across repeated runs, failures included, never a single successful run.

PAST FAILURES

We have seen AI programs fail before

The documented causes are operational rather than technical: no measured baseline, no agreed success threshold, governance retrofitted at security review, classification deferred until the design is already built. Three of those are closed by construction here - the baseline and the threshold are signed before work starts, governance is a property of the build, and classification is settled before the architecture rather than after it.

DEPENDENCY

This creates a new dependency on you

The measured gap at six, twelve and twenty-four months, an annual cold-start test with nothing of ours present, an export performed rather than documented, and a right to buy us out.

THE PERIMETER

Nothing real can leave our perimeter

Nothing real does. Structure crosses, code never crosses, personal data never enters. One audited exit, no fast path, no override.

CLASSIFICATION

De-identification is not declassification

We agree, and we said it first. The determination stays with your Business Data Executive.

PRODUCTIVITY

We have heard AI productivity claims before

Baselines scored on your own historical records, in the language they were written in, agreed before any target is discussed. An improvement target cannot be claimed without a measured baseline, and in most entities no such baseline exists.

COMPLIANCE

A provider's compliance posture is unproven until it is tested

It is, and you carry the consequence of it rather than us. Test it against sections [7](#), [8](#) and [10](#), which are written to be checked against published text rather than accepted.

The First Step

One working session. **You choose the task.**

We bring the System Brain our own practice runs on - available to interrogate, bounded, and in person. You pick the task, and we run it in the room.

Then we run the same task repeatedly and **show the distribution of outcomes, failures included**. A single successful run is a demonstration; a distribution is evidence. It is the only demonstration we think is worth anything, and it is the opposite of what a conventional vendor demonstration offers.

ONE SUCCESSFUL RUN



A DEMONSTRATION

THE SAME TASK, RUN AGAIN AND AGAIN



A DISTRIBUTION, FAILURES INCLUDED

ON A NAMED
INSTRUMENT

There is a policy basis for it, too: DGA's own instrument obliges entities to adopt proofs of concept, prototypes and sandboxes where possible, so this sits on a named instrument rather than on goodwill.

If the session is useful, a proposal follows. If it is not, you will have spent an hour and learned something checkable about what this class of system can and cannot do.

ARRANGE IT
WITH

Yasser Al Masri

AI TRANSFORMATION ARCHITECT

yasser@tutelex.ai · +966 54 444 7769

The seven words this document uses

Estate

An entity's technology taken as a whole: applications, databases, servers, networks, and the documents describing them. The **live** estate is what is serving users right now.

System Brain

A permanent, entity-owned model of how the estate works and why, which grades its own knowledge as verified, inferred, or an acknowledged gap.

Agentic

Software given a goal rather than a script: it plans, uses tools, checks its work, and acts.

Custody

The arrangement in which the entity holds every part of the platform, holds a record sufficient for its own engineers, and has its independence from us measured on a schedule rather than promised.

The Gate

The single audited exit from the perimeter, and therefore the only place that needs watching.

Frontier model

The most capable general-purpose AI models, all of which run outside any entity's perimeter. An **open-weight** model runs on your own hardware instead.

Blast radius

Everything a proposed change can reach. Stated before approval rather than discovered after it.

ON SOURCING

Every figure attributed to DGA is from DGA's own published reports. Where a figure is not DGA's, its source is named at the point of use. Every clause cited has been read in the issuing authority's own published text rather than in secondary commentary. We do not cite an unpublished internal target as though it were a national mandate.