

SHARQ

شرق

**Most AI wasn't built
for how the Gulf works.**


Yours can be.



The Sovereign
AI Issue

Edition 01
Company overview

Sep 2026
Dubai, UAE



Where general-purpose AI falls
short


The Arabic token premium,
measured

What a sovereign model includes

A pilot built around one workload




The Sovereign AI Issue



Why it matters

- 03**
Opener: where SHARQ fits
- 04**
AI that speaks your language
- 05**
Where general-purpose AI falls short
- 06**
The Arabic token premium



What we build

- 07**
What SHARQ builds
- 07**
Where it pays back first
- 08**
How it works: everything inside your boundary
- 09**
The compute is already here

PLUS: the method behind our token analysis, and the regulation it cites, are on pages 06 and 11.

Cover illustration: original render for SHARQ.



Working together

- 08**
Where SHARQ fits among your options
- 10**
The pilot
- 10**
The pilot scorecard
- 11**
Next step and contact
- 11**
Notes and sources

Where SHARQ fits

01

Why it matters. What SHARQ does, and why general-purpose AI costs Gulf organisations more than it should. Pages 04–06.

02

What we build. A specialist model you own, the UAE compute it runs on, and where it pays back first. Pages 07–09.

03

Working together. How SHARQ compares with your other options, and how a pilot proves the case. Pages 08, 10–11.

شارق



Feature
What SHARQ does

Words
SHARQ Artificial Intelligence

Illustration
Original render for SHARQ

AI that speaks your language

Specialist Arabic-English models, built on your own data and run inside your borders.

Most AI in use today was built for English first and trained on the open internet. For a bank in Dubai or a ministry in Abu Dhabi, that creates three problems at once. The model is weakest in the language customers actually use. It knows nothing about the organisation's products, policies or regulator. And it often runs somewhere the data is not allowed to go.

SHARQ closes those gaps with specialist models. We start from a strong open-weight model and post-train it on the organisation's own documents, workflows and terminology, in Arabic and English together. That includes Gulf dialect and the mid-sentence switches

between languages that real work involves.

The result is a model the client owns. It is tested on the client's real tasks before it goes live. It is trained and run on compute inside the country: the client's own servers, or sovereign cloud already operating in the UAE, such as an OCI Dedicated Region or Core42.

Two things we do not do. We do not build a foundation model from scratch, because strong open-weight models already exist and the value is in adapting them. And we do not ask anyone to take our word for results. Every engagement starts with one workload and a scorecard agreed in advance, so the case for going further is made on the client's own data. 



Fig. 1 Two linked rings, one cyan and one purple: Arabic and English trained as a single model.

We don't ask anyone to take our word for results. We agree the test first.

At a glance

Base: open-weight model
Languages: Arabic, English, Gulf dialect
Trained: on GPUs inside the UAE
Runs: your servers or UAE sovereign cloud
Weights: owned by the client
Start: one workload, scorecard agreed first

Section
The problem

Focus
Why generic AI falls short

Where general-purpose AI falls short

01

Language

Gulf teams mix Arabic, English and dialect in one sentence. Most models are tested one language at a time, in standard form, so the test rarely matches the work.

02

Location

Regulated data often has to stay in the country. In the UAE, banks must hold and store all consumer and transaction data within the UAE, which rules out many foreign-hosted AI services. Since February 2026 the Central Bank has run its own sovereign financial cloud.

03

Fit

A general model does not know your products, policies or regulator. Staff re-explain that context in every prompt, or correct the output afterwards, by hand.

A typical request

ممکن تشيک ال KYC status للعميل قبل ال onboarding؟

“Can you check the client’s KYC status before onboarding?” Gulf Arabic with English terms mixed in: ordinary at work, and hard for models tested in one language.

The cost shows up as extra compute, manual correction, and sensitive work that never reaches AI at all.

None of these gaps is a flaw in the models themselves. General-purpose systems are built to be good at everything for everyone. Gulf organisations need something narrower: a model that handles their mix of languages, knows their business and runs where their regulator expects.

The regulator is moving the same way. In February 2026 the Central Bank of the UAE launched a sovereign finan-

cial cloud for licensed financial institutions, built with Core42 as dedicated, isolated infrastructure. When the sector’s own regulator builds sovereign infrastructure, local models stop being a preference and become the baseline.

The first of the costs above is the easiest to measure, because it shows up in every request. The next page sets out what we found when we counted it.

All

consumer and transaction data held by UAE-licensed financial institutions must be stored within the UAE.

CBUAE Consumer Protection Standards, Art. 6, cl. 6.1.6.3

Sources


Central Bank of the UAE Rulebook, Consumer Protection Standards, Article 6: Protection of Consumer Data and Assets, clause 6.1.6.3.

CBUAE, sovereign financial cloud services developed with Core42, 25 February 2026.



The Arabic token premium

The same text takes more tokens in Arabic than in English. How many more depends on the model.



Method. Official Arabic and English texts of the UN Universal Declaration of Human Rights (NLTK udhr2 corpus), counted in full with each model’s published tokenizer, September 2026. **Caveats.** Token counts are not prices; per-token prices vary by provider. Formal Arabic only: Gulf dialect and mixed-language text were not measured. Independent research points the same way: Petrov et al. (NeurIPS 2023) measured a 4.40x premium for Standard Arabic on GPT-2 using FLORES-200.

5,251 → 2,378


Arabic tokens for the same text: GPT-4 tokenizer, then GPT-4o tokenizer. The model choice alone cuts the count by more than half.

Section
The product

Focus
Your specialist model

What SHARQ builds

An open-weight base model, post-trained around your organisation on compute inside the UAE, and delivered with everything needed to run it.



Bilingual

Trained on Arabic, English and Gulf dialect together, including requests that mix all three.

Specialist

Learns your products, policies and terminology from your own documents, so each request needs less instruction.

Sovereign

Runs on your premises or on sovereign cloud inside the UAE. You own the weights and can inspect the whole model.

Local training

Post-trained on GPUs inside the UAE, linked by RDMA networking. Your data is never exported to train it.

Delivered with

Evaluation on your own tasks

Post-training on UAE-hosted GPUs

Deployment, monitoring and scheduled refresh

Customer-facing AI is the front door. The savings usually sit behind it.

Start where the work is high-volume, narrow and regulated, and where the data should not pass through a general-purpose API. SHARQ sits alongside the assistants customers already see, not in place of them.

Banks and finance

KYC and compliance review

Credit memo drafting

Fraud alert triage

Policy and procedure search

Government

Citizen enquiries in Arabic and dialect

Policy and legal research

Bilingual correspondence

Case summaries for officers

Government • live now


The Abu Dhabi Judicial Department launched the first phase of its Judicial AI Platform in September 2026: case analysis and judgment drafting, with every output verified by a judge. A regulated workflow, running today.



Section
Working together

Focus
Process and options

Everything happens inside your boundary



🔴 Your data and the model weights stay inside this boundary, and inside the UAE, at every step.

Where SHARQ fits among your options

	General AI API	Consultancy build	In-house team	SHARQ
Where your data goes	Provider’s cloud; regions vary	Depends on the build	Your environment	Your environment or UAE sovereign cloud
Compute infrastructure	Foreign API clusters	Depends on the build	Hardware you buy and run	Local AI compute: GPU clusters inside the UAE
Arabic and dialect	General-purpose	Varies by team	Needs specialist hires	Post-trained on your bilingual data
Who owns the model	The provider	Varies by contract	You	You
How quality is proven	Public benchmarks	Acceptance tests	Tests you build	Evaluated on your own tasks
Ongoing effort	Low	Yours after handover	High: hiring and upkeep	Shared: we run and refresh it

Typical patterns, not claims about any named provider. Individual offerings vary.

The compute is already here

AI compute at national scale is being built inside the UAE. A sovereign model no longer has to trade local control for scale.

Sovereignty, as SHARQ defines it

Local compute + full model visibility

Keeping data in the country is only the start. A sovereign model is also trained and run on hardware inside the UAE, and nothing about it is hidden from you: you hold the weights, can inspect the model, and can trace how it was trained. Data localisation alone gives you none of that.

UAE compute, coming online

5 GW

UAE-US AI Campus, Abu Dhabi, at full build-out. The first 500 MW phase is due online by the end of 2026.

1 GW

Stargate UAE, a cluster on the same campus from G42, OpenAI, Oracle, NVIDIA, Cisco and SoftBank. First 200 MW due in 2026.

Where a sovereign model can run, inside the UAE

OCI Dedicated Region

Oracle cloud services inside Abu Dhabi's borders, on Core42 infrastructure with NVIDIA GPUs, so sensitive data does not leave the emirate.

Core42 sovereign cloud

G42's sovereign cloud and GPU compute, hosted in the UAE. Core42 also built the Central Bank's sovereign financial cloud.

Your own premises

Your servers, with the model packaged as a self-contained, self-hosted deployment that you control end to end.

How SHARQ post-trains

Post-training on local accelerators, linked by RDMA networking.

The compute provider for each engagement is agreed at scoping, from UAE-hosted options such as those above.

RDMA lets GPUs exchange data directly, which keeps training at enterprise speed. Because the accelerators sit inside the UAE, your data never has to be exported to train the model, and foreign data-export exposure is removed at the source.

Trained here. Run here.


Data, training compute and model weights stay inside the UAE at every step. Capacity figures are as announced by the operators; sources on page 11.



Section
Working together

Focus
The pilot

Start with one workload and a clear decision at the end



Durations and commercial terms are agreed at scoping. Security review and data-access approvals should start in Phase 1.

We'd rather show it on your data than put it on a slide.

Every pilot ends with a scorecard agreed before work starts. The figures come from your workload, measured against how you do it today. We do not quote savings in advance; we measure them.

The pilot scorecard

Task accuracy

Against your current process, on held-out work

Cost per 1,000 requests

Compute and licences at your expected volume

Correction rate

Share of outputs your team has to fix

Response time


Measured at your expected production load

Data boundary

Confirmation that data and weights stayed with you, inside the UAE

Next step
A scoping session

Pick one workload. We'll make the case on your data.



One scoping session with your business and technology leads. We choose the workload, agree the scorecard, and send a written pilot proposal. 📄

The founders

Ronit Dhansoia

Co-Founder

ronit@sarq.ai

Aman Kumar

Co-Founder

aman@sarq.ai

Suryavir Kapur

Co-Founder

suryavir@sarq.ai

SHARQ Artificial Intelligence · Dubai, United Arab Emirates

Notes

Results. SHARQ does not quote savings in advance. Pilot results are measured on each client's own workload, against a scorecard agreed before work starts.

Token analysis. Counts on page 06 come from the official Arabic and English texts of the UN Universal Declaration of Human Rights, tokenised in full with each model's published tokenizer in September 2026. Token counts are not prices. Formal Arabic only; dialect and mixed-language text were not measured.

Comparison. The options table on page 08 describes typical patterns, not any named provider.

UAE compute. Capacity figures on page 09 are planned and phased, as announced by the operators. Named platforms are deployment options, not SHARQ partners.

Images. Original renders made for SHARQ.

Sources

01 Central Bank of the UAE Rulebook, Consumer Protection Standards, Article 6: Protection of Consumer Data and Assets, clause 6.1.6.3. rulebook.centralbank.ae

02 Petrov et al., "Language Model Tokenizers Introduce Unfairness Between Languages," NeurIPS 2023. arxiv.org/abs/2305.15425

03 United Nations, Universal Declaration of Human Rights, official Arabic and English texts (NLTK udhr2 corpus).

04 SHARQ token analysis, September 2026: GPT-2, cl100k_base, o200k_base, Llama 3 and Gemma 3 tokenizers.

05 Central Bank of the UAE, "CBUAE Develops World's First Sovereign Financial Cloud Services," 25 February 2026. centralbank.ae

06 Gulf News, "Stargate UAE: first phase of 200MW to complete by 2026," 16 October 2025.

07 Zawya, "Largest AI campus outside of US to open in Abu Dhabi this year," 2026: first 500 MW phase of the 5 GW UAE-US AI Campus.

08 NVIDIA, "Oracle and NVIDIA Accelerate Sovereign AI, Enabling Abu Dhabi's AI-Native Government Transformation," 14 October 2025.

09 Nukta, "Abu Dhabi launches first phase of Judicial AI Platform," 28 September 2026.