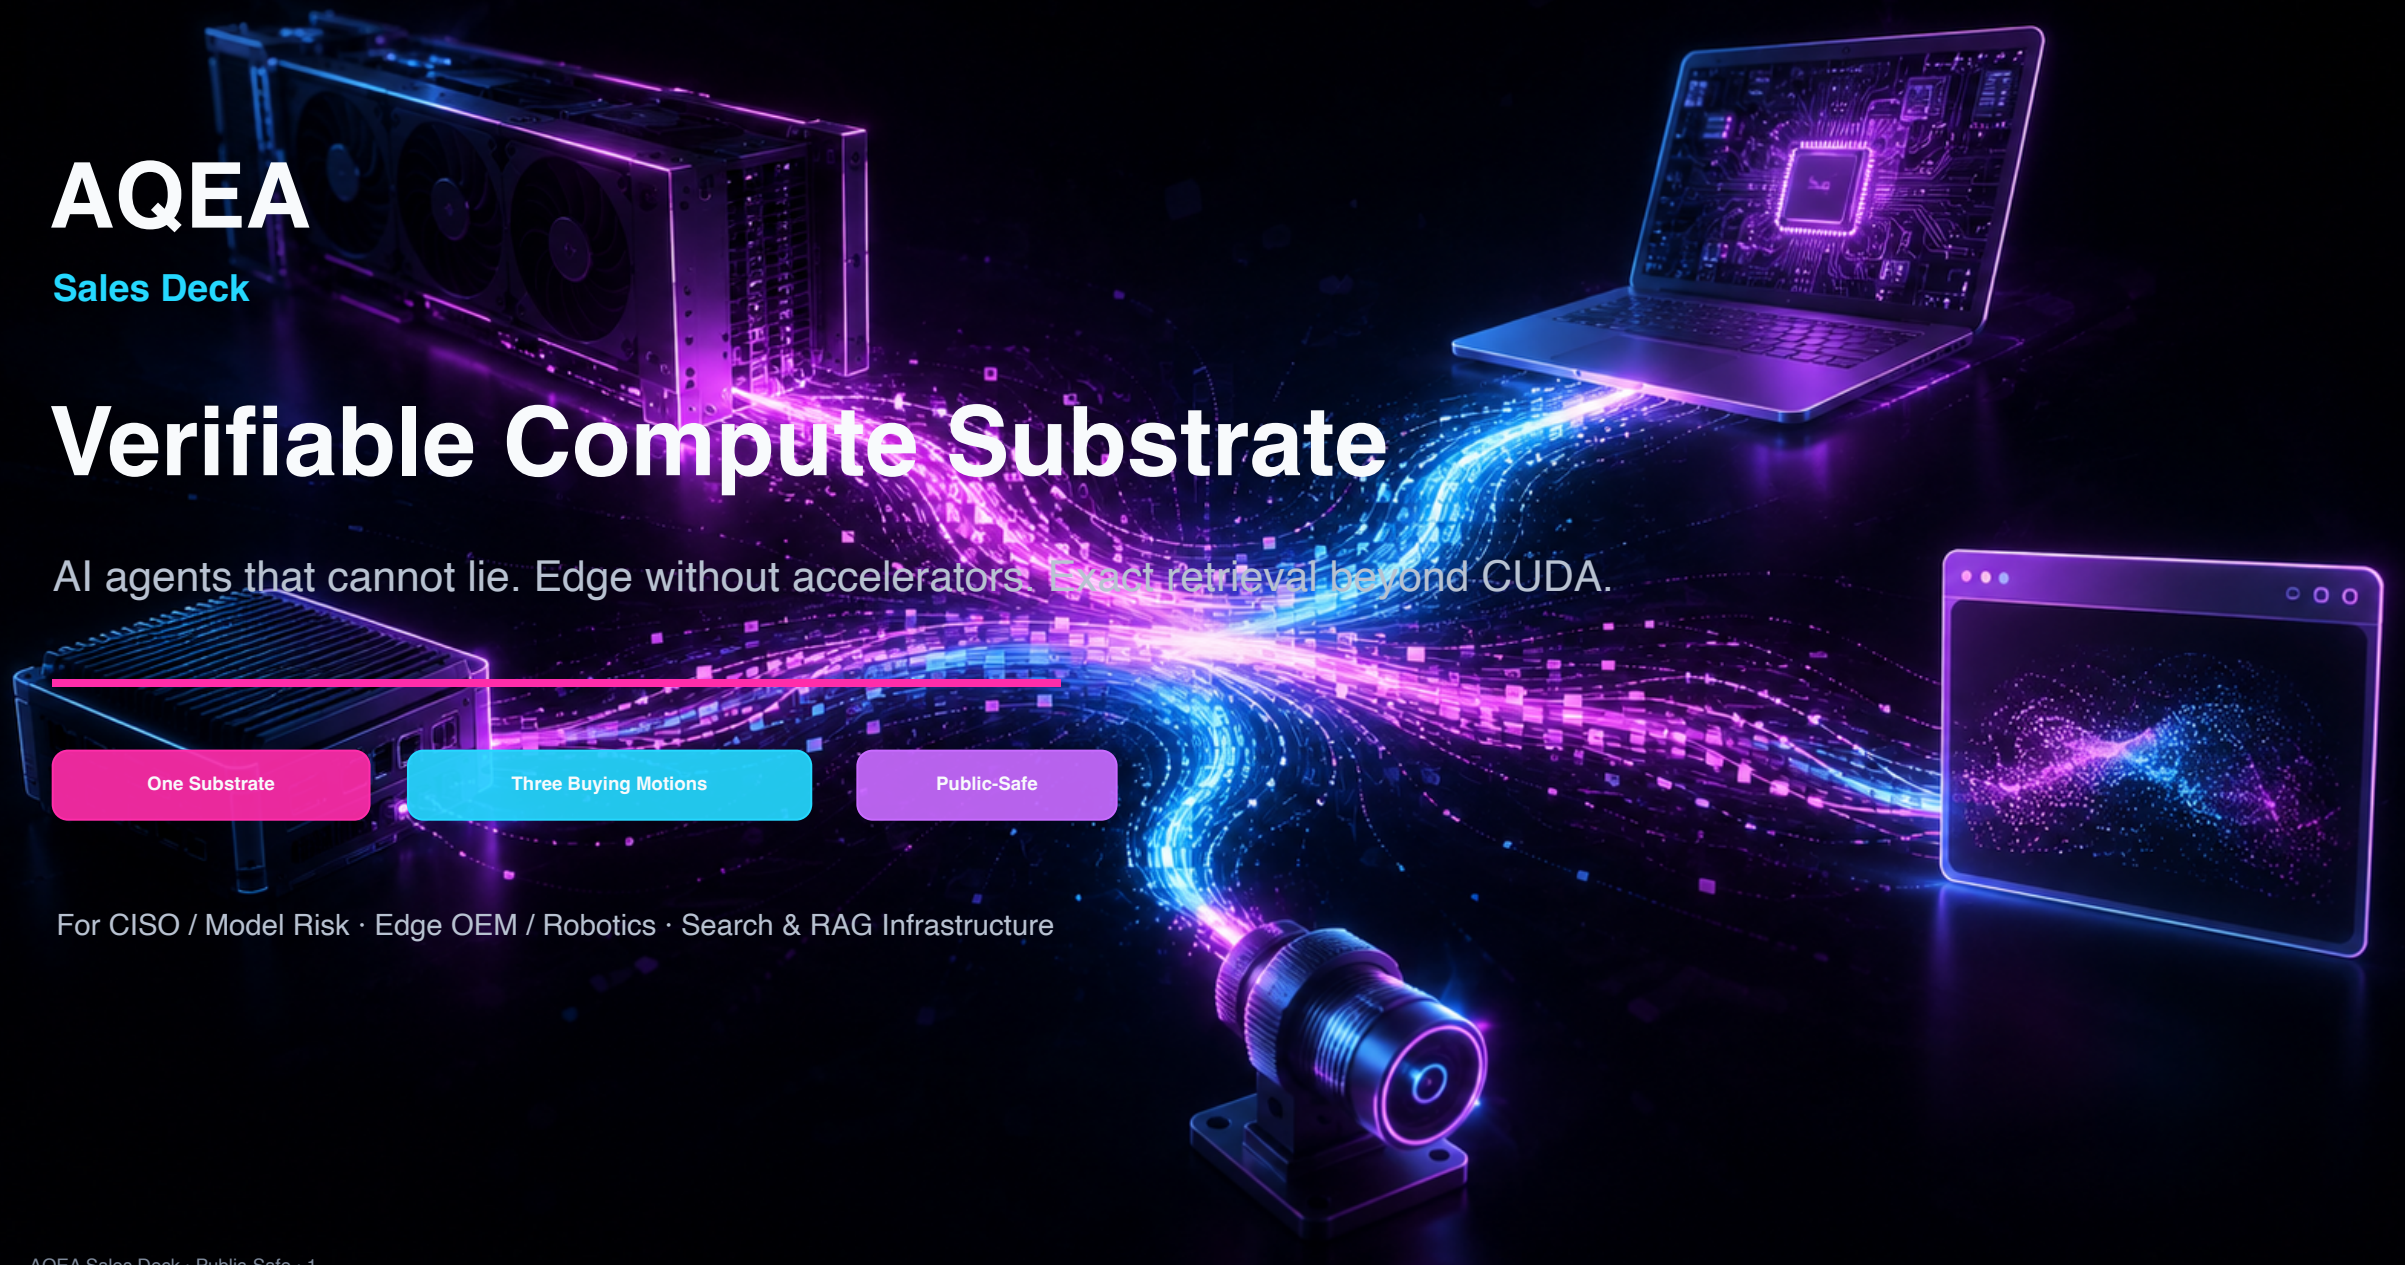

# AQEA

## Sales Deck

# Verifiable Compute Substrate

AI agents that cannot lie. Edge without accelerators. Exact retrieval beyond CUDA.

One Substrate | Three Buying Motions | Public-Safe

For CISO / Model Risk · Edge OEM / Robotics · Search & RAG Infrastructure

# The buyer problem is no longer “more AI”.

It is proof, cost, and hardware independence.

## Regulated agents

“Prove the agent did not do something forbidden.”

**CISO / Model Risk**

## Edge fleets

“Remove AI-chip BOM without losing recognition quality.”

**OEM / Robotics / Industrial**

## RAG search

“Exact answers without CUDA lock-in or 40GB indexes.”

**Search / Infra / Platform**

**AQEA turns these three pains into three separate entry points — backed by one protected substrate story.**

# One substrate. Three sales motions.

Start with the customer's pain. Land one wedge. Expand across the platform.

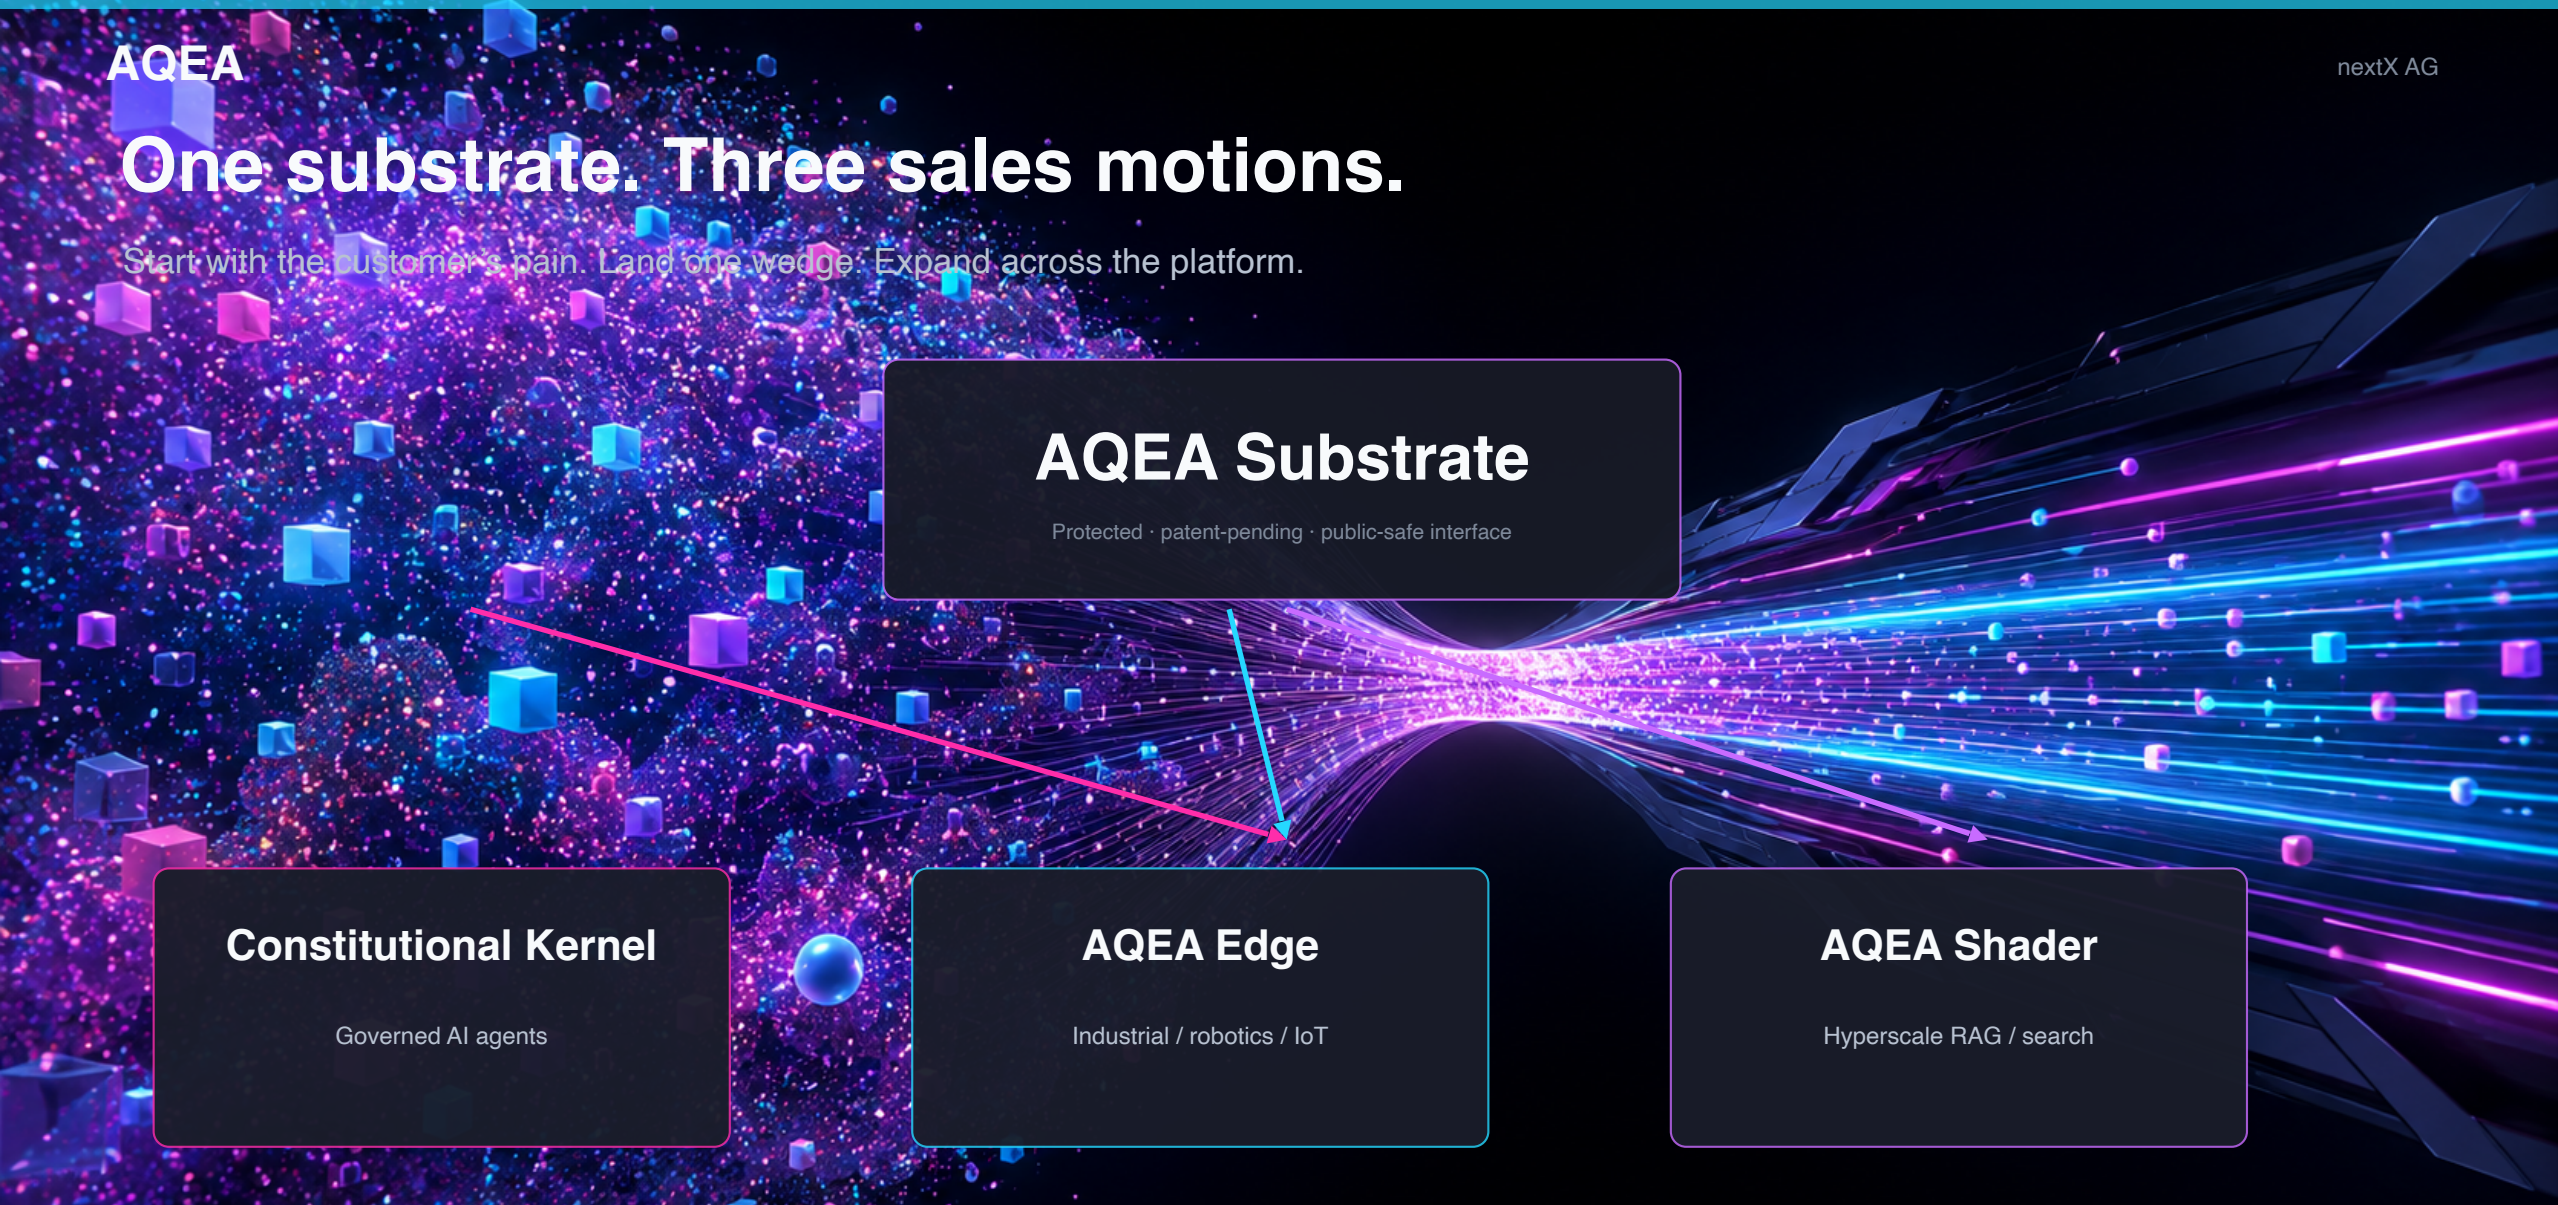


**Sales framing: sell the outcome first. Keep mechanism under NDA.**

# Who buys AQEA?

Three budget lines, one expansion path.

| Buyer | Entry wedge | Business pain | Next step |
|---|---|---|---|
| **CISO / Model Risk** | Agent governance | Compliance risk, audit evidence, human approval | Design partner / governance pilot |
| **Edge OEM / Robotics** | Accelerator removal | BOM, energy, vendor lock-in, fleet validation | Joint benchmark / edge pilot |
| **Search / RAG Infra** | Exact cross-vendor retrieval | Index size, latency, CUDA concentration | NDA engineering evaluation |

# Wedge 1: Constitutional Kernel

AI agents that cannot lie about what they did.

## Customer pain

- LLM agents can write code, call tools, mutate systems
- Audit trails written by the agent are not evidence
- Regulated teams need deterministic replay and approval gates

## What AQEA does

- Rule-Core sealed at build time
- Five-outcome action gate
- Hash-chained WAL + Ed25519 approvals
- Sandbox-first tree-hash invariant

**184/184**
deterministic tests

**~14s**
end-to-end demo

**5**
exhaustive outcomes

**20**
refusal paths tested

**Positioning: not a wrapper, not an LLM, not a policy PDF — a governance substrate between the agent and the system.**

# Wedge 2: AQEA Edge

Similarity recognition without a dedicated AI accelerator.

## Replace this cost line

- Predictive maintenance
- Robot fault detection
- Quality-control fingerprints
- Wearable / hardware anomaly
- Wake-word / acoustic events

**13/13**

domains above floor

**174%**

fall-prediction recall vs Float

**9.4x**

energy advantage peak

**$0**

new AI-chip BOM

## Sales claim

If the workload is similarity-classification, AQEA can move the recognition path from vendor-locked AI silicon to CPU SIMD or shader backends already inside the device.

**Best fit: OEMs with fleets, BOM pressure, and long product lifecycles.**

# Wedge 3: AQEA Shader

Exact retrieval beyond CUDA.

**1.85x**

faster than Float-FAISS-GPU at 10M

**23x**

smaller index storage

**100%**

Top-K bit-identity to brute force

**4/4**

GPU backends bit-identical

## Why it sells

- Cuts memory traffic
- Preserves exact recall
- Avoids CUDA-only dependency
- Works toward browser / WebGPU use cases

**Best fit: search infra, RAG platforms, hyperscalers, sovereign cloud, hardware vendors.**

# Where the ROI comes from

AQEA is sold as avoided cost + reduced risk + hardware freedom.

## Risk avoided

Agent actions become replayable, approval-bound, and policy-precedence-rigorous

CISO / Model Risk

## BOM removed

Similarity-classification can run without per-device AI accelerator silicon

Edge OEM

## OpEx reduced

Smaller indexes and lower memory traffic reduce storage, energy, and GPU dependency

Search infra

## Strategic freedom

One encoded artifact can move across vendor targets without answer drift

Platform teams

# The first 30 days should feel easy

A sales motion built around proof, not claims.

## 1 Discovery

Identify wedge, buyer, current baseline, success metric

## 2 NDA + data scope

Partner shares corpus / sensor set / governance workflow

## 3 Joint benchmark

Run AQEA against their baseline on their hardware

## 4 Decision packet

Bench report + integration plan + commercial proposal

# Packages that make buying simple

Three options, increasing commitment.

## Engineering Evaluation

**4–6 weeks**

Fastest route to proof

Reproducible benchmark on partner workload + hardware

## Integration Pilot

**8–12 weeks**

Best route to budget

Shadow-mode integration in non-production path

## Strategic Co-Development

**Custom**

Best route to exit / alliance

Hardware tuning, vertical encoder, or bespoke deployment

**Default CTA: “Let us run the first benchmark on your own workload.”**

# Objection handling

Keep the answers sharp. No mechanism leakage.

| | |
|---|---|
| **“Is this just a wrapper?”** | No. Kernel sits between the agent and the system; outcomes are deterministic and audit-chained. |
| **“Does Edge replace all AI inference?”** | No. It replaces similarity-classification paths, not generation or segmentation. |
| **“Is Shader approximate search?”** | No. It is positioned as exact Top-K bit-identity to brute force on the encoded corpus. |
| **“Do we need to change encoders?”** | No. Bring your own encoder; AQEA consumes deterministic encoder output. |

# Choose the entry wedge

The sales call should end with one concrete next step.

## Governed agents

**Book a Constitutional Kernel workshop**

Bring: restricted actions, approval rules, audit requirements

## Edge fleet

**Book a sensor benchmark**

Bring: encoder output, reference set, target hardware

## RAG / search

**Book a retrieval benchmark**

Bring: corpus, queries, baseline, hardware target

# AQEA

## Verifiable Compute Substrate

Sales CTA: prove it on the buyer's workload.

**Capital + strategic pilots + hardware allies**

**nextX AG**

s.karim@nextx.ch

**NDA benchmark scope returned within five working days**