AQEA Gen-5 Embedding

Structured information-geometry substrate for embeddings that need to be smaller, deterministic, cross-vendor, reversible, and exact where it matters.

NOT A NEW FOUNDATION MODEL

GEN-4 IN · GEN-5 OUT

PATENT-PENDING

What changes vs Gen-4?

Gen-4 produces one dense float vector. AQEA projects encoder output into a structured substrate with internal channels, deterministic encoding, compressed storage, ranking preservation, and reversible decode modes.

23x

smaller 10M index vs FP32 float baseline

1.85x

faster than Float-FAISS-GPU at 10M

100%

Top-K bit-identity verified vs brute force

4/4

GPU backends bit-identical

13/13

domain x encoder validations above floor

174%

sensor-domain recall vs Float-FAISS

Gen-4 encoder output → **AQEA Gen-5 substrate** → **CPU / GPU / WebGPU path** → **Exact search · edge recognition · audit decode**

What Gen-5 Embedding gives you

A product-facing view of the AQEA substrate capabilities - public-safe, benchmark-led, and designed for partner evaluation.

Structured channels

Multi-channel organization; operations can be separated instead of collapsed into one float axis.

Deterministic encoding

Byte-identical outputs across heterogeneous CPU and shader targets.

Structural compression

23-29x smaller representation class while preserving nearest-neighbor ordering.

Exact ranking mode

Top-K equality to brute-force reference on the encoded corpus: same documents, order, and distances.

Noise-resistant sensors

On raw signal domains, encoded ranking can exceed float baseline by filtering per-bin noise.

Reversible decode

Deployment-time modes for audit-fidelity, general-purpose, or pure-retrieval operation.

Public benchmark anchors

10M retrieval latency	6.27 ms p50 vs 11.60 ms baseline
Index footprint	1.55 GB vs 36.0 GB at 10M
Reversible decoding	99.68% retrieval-equivalent at 1M
Energy	1.73x more efficient at 1M; Edge up to 9.4x lower

Best-fit deployment surfaces

Hyperscale RAG / semantic search	exact retrieval beyond CUDA
Industrial / robotics / IoT	similarity recognition without AI accelerator
Regulated audit chains	decode / verify evidence without re-acquisition
Cross-vendor infrastructure	encode once; deploy across targets

Integration model Bring your encoder. AQEA consumes deterministic encoder output, creates a compact Gen-5 representation, and runs retrieval / recognition / verification on CPU SIMD or standard shader backends.

CTA: 4-6 week evaluation under NDA

Implementation details disclosed under NDA only

Security, Audit & Compliance Layer

Gen-5 is not a black box. It is verifiable, immutable where it matters, and audit-native by design.

CONTROL PLANE

Immutable by construction

A build-time-sealed Rule-Core gates proposed actions before materialization.

Forbidden targets resolve to SUPPRESS_POLICY_DENY with highest precedence.

Positive evidence cannot outvote a policy denial: tests passing, performance gains, or signatures do not resurrect a refused action.

5 OUTCOMES | 184 TESTS | 20 PATHS

Verifiable chain of custody

Every action becomes evidence: gated, signed, sandboxed, and replayable.

Proposal	Action Gate	WAL	Signature	Sandbox	Verified
LLM / encoder output is only a proposal source	five deterministic outcomes	pre-effect SHA-256 chain	Ed25519 approval binding	tree-hash invariant	replayable audit record

AUDIT

Hash-chained WAL

Pre-effect SHA-256 chain. Post-hoc edits break downstream SHA.

AUTHORITY

Ed25519 approval

Binds chain, patch, target, scope and expiry. Local verification.

SAFETY

Sandbox-first

Live tree unchanged until approval. Tree-hash asserts invariant.

COMPLIANCE MAP EU AI Act transparency / oversight / robustness · Banking model risk · 21 CFR Part 11 · HIPAA / GDPR auditability