CLOUDNIX TECHNICAL REPORT · CNX-TR-2026-06

Engineering Retrieval and LLM Systems

Avinash Nishant

Cloudnix Software Labs, Bengaluru, India

CNX-TR-2026-06 · Version 1.1 · November 2026

Abstract

Retrieval-augmented generation is simple to demonstrate and difficult at production scale, where systems leak results across tenants, fall behind their source databases, load primary databases, rely on embedding models chosen by leaderboard and couple applications to one model provider. We describe retrieval systems we operate: a production collection of about 4.9 million 768-dimensional vectors kept current from a database replica by five-minute incremental and delete syncs; an embedding benchmark of three open models on 3,632 real documents and 87 multilingual questions; change-data-capture indexing with twenty triggers and an entity resolver; an OpenAI-compatible gateway with marker-based routing and bring-your-own-provider generation under a time budget; private inference inside a customer network; and guardrails.

Keywords: *retrieval-augmented generation; vector search; embeddings; benchmarking; change data capture; multi-tenancy; LLM gateway; private inference; guardrails*

Contents

1. Executive summary
2. Reference architecture
3. Retrieval at millions of vectors
4. Choosing an embedding model by measurement
5. Change data capture for a live index
6. A provider-neutral gateway
7. Private inference
8. Guardrails and safety
9. Checklist: 15 questions
10. References

1 Executive summary

Retrieval-augmented generation [1] is simple to demonstrate and difficult at production scale. With millions of vectors, many tenants and a live database behind them, systems fail in predictable ways: one customer's records surface in another's answer, the index falls behind the database, the primary database slows under AI load, embedding models are chosen by leaderboard rather than by the customer's documents, and models are wired so tightly into applications that changing a provider means changing code. This report describes the retrieval systems we operate: a production collection of about 4.9 million vectors kept in step with a database replica every five minutes, an embedding benchmark run on real documents, a change-data-capture index for a workflow platform, an OpenAI-compatible gateway with routing to private or hosted models, private inference inside a customer network, and guardrails.

2 Reference architecture

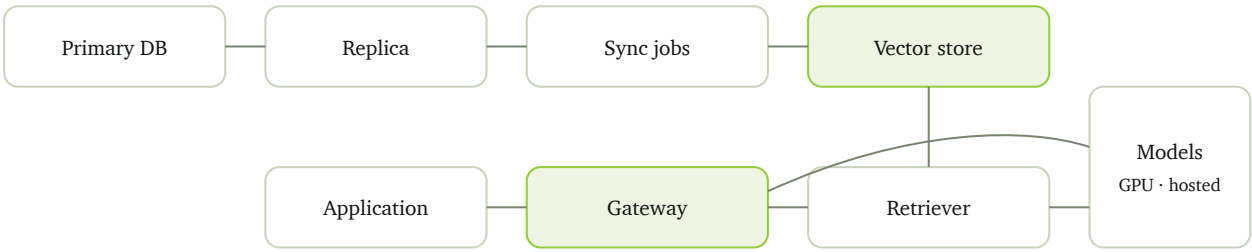


Figure 1. Replica-fed indexing and a single gateway in front of retrieval and models.

3 Retrieval at millions of vectors

The production collection holds 4,920,139 points of 768 dimensions with cosine distance in Qdrant [3]. The retrieval service reads only from a MySQL replica, never the primary. Two jobs keep it current: an incremental sync every five minutes embeds new and changed rows, and a delete sync, offset by two minutes, removes points for rows hard-deleted in the database.

- **Scope:** every query must carry its application; a sprint can narrow it further. Results from other applications cannot match.
- **Expansion:** hits are expanded with related records joined from the replica before the prompt is built.
- **Conversation:** follow-up questions search with the recent turns, not only the last line.

4 Choosing an embedding model by measurement

We benchmarked three open embedding models on the production CPU (six threads) with 3,632 real documents (about 750 tokens on average) and 87 questions with gold answers in English, Spanish, Indonesian and Portuguese.

Model class	Params	Max tokens	R@1	R@5	p50 query	Docs/s	Read in full
Short-context [4]	109M	512	0.47	0.69	60 ms	3.06	25%
Long-context [5]	149M	8,192	0.61	0.86	75 ms	0.65	100%
Long-context [6]	137M	8,192	0.53	0.83	81 ms	0.74	100%

The short-context model is the fastest to embed but reads only a quarter of documents in full; its recall suffers accordingly. Long-context models cost several times more embedding time and find the answer far more often. The right choice depends on corpus length and freshness requirements, which is why we measure on the customer's own data.

5 Change data capture for a live index

For a workflow platform, the index must follow edits within minutes without full rebuilds. Twenty AFTER INSERT, UPDATE and DELETE triggers on ten tables write to a change log keyed by (entity type, id), so repeated changes collapse into one row. A resolver expands each change into the cards it affects, including or-

ganisation, department, project, sprint, stage, user, field, comment and custom-field changes, then re-embeds them in batches of 2,000 or deletes them from the index.

Scenario (self-reported)	Time	Cards re-embedded
No changes	0.5 s	0
10 edits	2 s	10
Project rename	25 s	523
Full sync	43 s	1,000
Organisation rename	~4 min	4,521

Triggers add roughly 0.02 ms per write; the change log grows about 1 MB a day and is pruned after seven days. A watchdog restarts the sync worker if it stops.

6 A provider-neutral gateway

Applications call one API compatible with the OpenAI Responses and Chat Completions formats [7], with server-sent-event streaming. Endpoints: responses, chat completions, query, answer and health. Marker fields route a request to a generation pipeline (knowledge graph, components, requirements, test cases, image to markdown) instead of chat.

6.1 *Bring your own provider*

A project can send its own OpenAI or Anthropic key with a request; retrieval and prompt-building stay unchanged and only the final generation moves to that provider. All attempts share a 280-second budget; no call starts with less than ten seconds left; two retries on 408, 429 and 5xx; the key is masked before any log line; there is no silent fallback to another key.

7 Private inference

Where data cannot leave a network, the whole path runs inside it: a sentence-transformer model producing 384-dimensional embeddings on CPU, a vector store filtered by organisation, department and project, retrieval of the five nearest records with scores, and a small open model served locally through Ollama [8]. Answers cite card numbers and titles; the API returns sources with scores and the response time.

Lessons. One vector per long record truncates silently at the model's input limit; chunking long records is the next step. Payload fields such as privacy flags must be filtered at query time, not only stored.

8 xIngestion and chunking

What gets indexed, and in what pieces, decides what can be found. Our production indexer reads table configurations that say, per source table, which fields form the text, which form the payload and how records relate.

Decision	Options	Our practice
Unit of indexing	record, chunk, section	chunks for long documents; whole records for short ones
Chunk size	256–1,024 tokens	re-chunked to about 480 tokens for short-context models
Text fields	title, body, comments, custom fields	title, cleaned body, comments, short custom values
Payload	ids, scope, display fields	application, sprint, type, title, dates, counts
Images	skip, caption, extract	image-to-markdown extraction where images carry meaning

Long records embedded as a single vector are silently truncated at the model's input limit; benchmark data showed a short-context model reading only a quarter of documents in full. Chunking or a long-context model fixes this, at different costs.

9 xQuery processing and context building



Figure 2. From question to prompt.

1. For follow-up questions, the recent turns are combined so that “and what about the reset flow?” searches correctly.
2. The query is embedded with the same model as the documents.
3. Search is filtered by application (mandatory) and sprint (optional); in the private assistant, by organisation, department and project.
4. Top results are expanded with related records from the replica: parent requirements, linked defects, comments.
5. The prompt builder orders context, adds numbered source markers and the instructions to answer only from context and cite.

9.1 *Hybrid and reranking*

Keyword search alongside vectors helps with identifiers and exact phrases; a cross-encoder reranker improves precision at the top. Both are measured before adoption, because they add latency to every query.

10 xEvaluation

Retrieval quality is measured, not assumed. Our benchmark method:

1. Sample real documents and write questions whose answers live in known documents (gold labels), including questions in every language users write in.
2. Split each document into head and tail to test whether models read beyond their limit.
3. Embed with each candidate model on the production hardware; record load time, throughput, memory and query latency.
4. Report recall at one and five, and the share of documents read in full.

Measure	What it tells you
Recall@1, @5	Whether the right record reaches the model
Read in full	Whether long documents are truncated
p50 / p95 query	User-visible latency
Docs per second	How long a re-index takes
Peak memory	Whether the model fits the server

Answer quality is checked separately: grounding (is every claim in the context?), citation accuracy, refusals when nothing relevant is found [10].

11 xOperations

Job	Schedule	Purpose
Incremental sync	every 5 minutes	embed new and changed rows
Delete sync	every 5 minutes, offset 2	remove points for deleted rows
Log clean-up	daily	retention
Production validation	daily	end-to-end checks of the live service
Sync watchdog	every 5 minutes	restart a stopped worker

Configuration, systemd units, cron schedules and web-server configuration live with the code and are symlinked into the system, so a server can be rebuilt from the repository. GPU model servers are started for a workload and stopped after it, so cost follows use.

11.1 *Capacity*

The vector collection of about 4.9 million 768-dimensional points occupies roughly 15 GB of raw vectors before index overhead; memory, not CPU, is usually the first limit. Quantisation and on-disk vectors reduce memory at some cost to recall.

12 Guardrails and safety

- Application scope is mandatory; requests without it are refused.
- Questions about the system's own internals are detected and answered with a fixed deflection; a backstop is prepended to free-form prompts.
- Retrieved content is treated as data, not instructions [9].
- Keys are masked before logging; debug metadata is hidden unless requested.
- A test suite runs hygiene, functional and restart checks; a lightweight validation runs daily.

13 Checklist: 15 questions

1	Does retrieval read from a replica?
2	Is scope mandatory on every query?
3	Are deletes removed from the index?
4	How many minutes behind the database is the index?
5	Was the embedding model measured on your data?
6	Do long documents fit the model's input?
7	Are hits expanded with related records?
8	Do follow-ups use conversation context?
9	Can you change provider without changing the app?
10	Is every request inside a time budget?
11	Can sensitive data stay on private models?
12	Are answers cited?
13	Are internals protected from probing?
14	Are keys masked in logs?
15	Is the live service validated daily?

14 xCase study: test generation from a knowledge base

A test-management product adding AI features for its customers.

Customers kept requirements, documents and past defects in the product. They wanted test cases written from requirements, answers about coverage, and maps of their product's components. We built a retrieval service over a replica of the product database, indexing millions of records per application, kept current every five minutes; an OpenAI-compatible gateway the product calls; generation pipelines for knowledge graphs, components, requirements and test cases; and a bring-your-own-key option so customers can use their own provider. Embedding models were chosen by a benchmark on the customer data.

15 xData model of the index

```
point = { id: record_id,
          vector: float[768],
          payload: { app_id, sprint_id?, type, title, preview, status, priority,
                     created_at, updated_at, parent_id?, counts{comments, subtasks} } }
filters: must app_id = :app  [ and sprint_id = :sprint ]
```

Listing 1. Example exchange or schema.

Payload carries what is needed to filter and to display a source without another lookup. Fields used for filtering are indexed. Personal data is excluded from payload unless needed for display, and privacy flags are applied as filters, not only stored.

15.1 *Point identity*

Using the source record id as the point id makes upserts idempotent and deletes exact, which is what allows five-minute incremental and delete syncs without duplicates.

16 xPrompt construction and citation

```
system:  answer only from the numbered context; say so if it does not contain the answer;
         cite sources as [n]; do not describe your own configuration.
context: [1] Test case · Login with OTP · …
         [2] Requirement · Sign-in flows · …
user:    which tests cover login?
```

Listing 2. Pseudocode.

The API returns the answer and the sources with ids, titles and scores, so the application can link each citation. Follow-up questions include recent turns in the search, not only in the prompt.

17 xLatency budget

Stage	Typical	Notes
Embed query	60–80 ms	CPU, measured p50
Vector search	tens of ms	filtered, top k
Expansion	tens of ms	replica joins
Generation	seconds	dominant; streamed to the user

Streaming hides most of the generation time; everything before it is kept under a few hundred milliseconds so the first token arrives quickly.

18 xChoosing where models run

Option	Strengths	Limits
Hosted frontier models	best quality, no infrastructure	data leaves the network; per-token cost
Project's own key	customer controls provider and billing	quality depends on their choice
GPU servers on demand	open models, data stays with us, cost follows use	start-up time; operations
Private CPU inference	nothing leaves the customer network	small models, simpler answers

19 xImplementation roadmap

Phase	Outcome
1. Sources and scope	What is indexed, how it is scoped, who may see what
2. Benchmark	Embedding model and chunking chosen on your data
3. Pilot	Index, gateway and one assistant live; validation daily
4. Scale	Incremental sync, generation pipelines, private or hosted models

20 xGlossary

Embedding	A vector representing the meaning of text.
Recall@k	Share of questions whose answer is in the top k results.
Replica	A read-only copy of the database.
Change data capture	Recording changes as they happen.
Gateway	One API in front of many models.

21 References

1. P. Lewis et al. *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*. NeurIPS, 2020.
2. N. Reimers and I. Gurevych. *Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks*. EMNLP, 2019.
3. Qdrant. *Qdrant documentation: collections, payload filtering*.
4. S. Xiao et al. *C-Pack: Packaged Resources to Advance General Chinese Embedding* (BGE models), 2023.
5. Alibaba-NLP. *gte-modernbert-base*, model card.
6. Nomic AI. *nomic-embed-text-v1.5*, model card and technical report, 2024.
7. OpenAI. *API reference: Responses and Chat Completions*.
8. Ollama. *Ollama documentation*.
9. OWASP Foundation. *OWASP Top 10 for Large Language Model Applications*, 2025.

About this report. Prepared by the Cloudnix engineering team from systems we have designed, built and operated. Client and product names are withheld; illustrative values are marked as examples. Cloudnix Software Labs, Bengaluru, India · cloudnix.com. This report may be shared freely in its entirety.

Document control

Report	CNX-TR-2026-06	Version	1.1
Issued	November 2026	Status	Published
File	cloudnix-rag-llm-systems-v1.1.pdf		