

AX Radar Was Built Exclusively For Humans, Leaving AI In The Dark



The Original State (Human Face)

- Accessed via `news.ax0x.ai`
- Rich web dashboard for curating feeds and reading editorial commentary.

The Blind Spot (Agent Reality)

- Claude has no access to the 6,821 enriched items or 59 sources.
- Result: Requesting a briefing yields an AI shrug; it requires manual screenshots to parse context.

The Fix

- Build a second face allowing agents to read, search, and mutate the radar directly, with zero glue code.

Treating API Tokens Like Passwords Is A Recipe For Database Meltdown



The Wrong Reflex (Bcrypt)

- **Design:** Built for low-entropy human passwords (30–40 bits).
- **Mechanism:** Deliberately slow to prevent offline brute-force.
- **Query Speed:** $O(n)$. Requires a full-table scan with per-row comparison.
- **Outcome:** The moment QPS spikes, the database melts.

The Right Tool (SHA-256)

- **Design:** Built for Bearer tokens (32 random bytes = 256 bits of entropy).
- **Mechanism:** Fast hashing. You cannot brute-force the heat death of the universe.
- **Query Speed:** $O(\log n)$. `SELECT * FROM api_tokens WHERE hash = $1`. Exactly one B-tree hop.
- **Outcome:** Safe, scalable authentication that survives traffic spikes.

MCP Provides The Pipe, But The Agent Still Needs A Manual



The Pipe: MCP

- **What it does:** Exposes tool signatures automatically (name, params, return shape) to Claude Desktop or Cursor.
- **The Gap:** Returns raw scores (e.g., importance: 72) but provides zero semantics on what that score actually means.

The Manual: Claude Code Skills

- **What it does:** A .md file in ~/.claude/skills that triggers on phrases like 'brief me on the radar'.
- **Domain Knowledge Injected:**
 - * Defines HKR axes and Featured vs. P1 thresholds.
 - * Dictates Semantic vs. Lexical search (gist vs. exact-string).
 - * Enforces guardrails: Never save speculatively without explicit intent, and YouTube sources are never excluded.

Semantic Search Shipped In 24 Lines Of SQL By Riding Existing Rails

No Vector DB Needed:

Skipped Pinecone and custom tuning entirely.

The Pre-laid Tracks:

text-embedding-3-large vectors in halfvec(3072) and HNSW indexes already existed for cross-source dedup clustering.

The Express Route:

Exposing semantic search only required embedding the query and running ORDER BY embedding <#> \$q (24 lines of SQL).

The Stats:

- Cost: \$0.00002 per query.
- Latency: ~250ms p50 end-to-end.

DATABASE:
HNSW INDEX

SEARCH

Modern Products Require A Single Truth Served Through Multiple Faces



The Single Backend: Both faces share the exact same Postgres query layer. The backend stays identical. Building the back-half of a product means adding access modes for non-human consumers, not duplicating logic.