



The most cost-effective way to access Azure OpenAI with no capacity limits.

CogCache is the most cost-effective way to access Azure OpenAI with no capacity limits.

By leveraging a global network of Azure supercomputing clusters, CogCache provides the highest performance and lowest cost on the market for Azure OpenAI tokens, enabling companies to unlock the full potential of generative AI.

With CogCache, you can now:

- Benefit from the lowest cost on the market for Azure OpenAI tokens
- Achieve the highest performance with the fastest response times, lowest latency and full LLM grounding and alignment
- Access the most advanced LLMs with no capacity limits
- Pay only for what you need, with monthly and annual subscriptions
- Get started in minutes with a simple purchase process



Azure OpenAI

UP TO

50%

CHEAPER

UP TO

100x

FASTER

COMPLETE

AI control AND alignment

Reduce Costs and Carbon Footprint

Save up to 50% on LLM costs with CogCache for similar, reoccurring prompts and cut your carbon footprint, making your AI operations more sustainable and cost-effective.

Asynchronous Alignment Scoring

CogCache uses Dynamic Constitutional AI to analyze and score cache content asynchronously, ensuring alignment and grounding of responses to uphold your brand integrity and comply with governance requirements.

Full-stack LLM Observability

Gain real-time insights, track performance key metrics and view all the logged requests for easy debugging.

Fast Cache Retrieval

100x your LLM response times with our two-tiered system uses high-speed in memory hashing and vector search to radically reduce response times and lower token rendering costs.

Self-Healing Cache

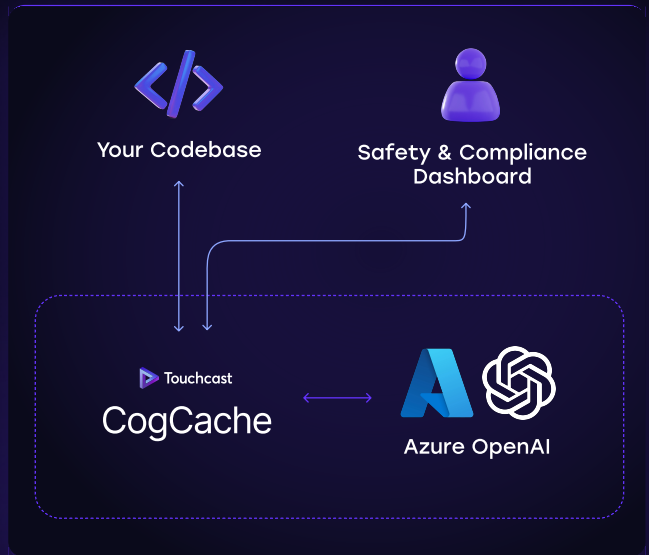
CogCache's self-healing mechanism identifies and mitigates misaligned responses, permitting automatic and human-approved adjustments for cache integrity and system reliability.

Temporal Relevance Tracking

The cognitive cache has the ability to discern if a cached result is temporally relevant and might no longer be applicable based on its content.

Your CogCache dashboard allows teams to monitor generative flow, set policy, and audit content with confidence.

Complete transparency and auditability of all generated content



Get started in minutes: Simply point your Azure OpenAI calls to your CogCache instance. No code changes needed.

ITEM	ALIGNMENT	HTS	DATE
sws/POC dev: What is the metaverse?	93	2	01/02/2023
sws/POC dev: How can we integrate it into our sales process?	93	1	01/02/2023
sws/POC dev: Why banks should be more oriented to Gen Z?	55	1	01/02/2023
sws/POC dev: What will be the impact of 5G on Fintech and what will be the economic impact on the overall market?	80	1	01/02/2023
sws/POC dev: How can we win more business using Gen AI?	65	13	01/02/2023
HR Knowledge Base: Show me our policies on data retention	40	20	01/02/2023
sws/POC dev: Provide me a list of best practices within Microsoft for increasing brand awareness	90	5	01/02/2023
Employee Support Center: What have you done to avoid bias or discrimination in your responses?	85	2	01/02/2023

Ability to edit, align and correct cache entries

Always Available

Ensure throughput is always available at low latency with our reserved processing capacity (whether used or not).

Flexible

Enjoy complete flexibility with no monthly commitments required.

Smart

Empower your teams with the ability to monitor generated content, ensuring observability and control over AI processes.

Cost-effective

Minimize costs by eliminating the need to use tokens on previously generated content.

Easy

No implementation resources required - simply switch your endpoints and get started in minutes.