

DropTicks AI Article

Architecture Guide: self-hosted inference
LLM Infrastructure | Mar 09, 2026

A system-design view for planning production implementation. how to serve open models with throughput, batching, and cost control

Article

Architecture Guide for self-hosted inference

This article explains how to serve open models with throughput, batching, and cost control. It is written for teams that need production behavior, not only a convincing prototype.

Start by defining the workflow boundary. List the inputs, outputs, permissions, failure states, and human review points. Most software and AI failures become easier to manage when the boundary is explicit.

The implementation should separate product intent from infrastructure mechanics. Keep validation, storage, observability, retries, and deployment rules outside of prompts or UI copy. Those rules belong in code, configuration, and tests.

A strong delivery plan includes:

- â€¢ A small reference workflow that proves the approach.

- â€¢ Test data that represents real edge cases.

- â€¢ Logging and metrics that expose failures instead of hiding them.

- â€¢ A rollback path for bad releases.

- â€¢ Documentation that explains ownership and maintenance.

For DropTicks projects, the best solution is usually the one that makes the system easier to inspect. If a team cannot explain why a tool was selected, how it fails, and how it is measured, the system is not ready.

Related source links:

- â€¢ <https://docs.vllm.ai/>

- â€¢ <https://huggingface.co/docs/text-generation-inference/index>

Web link

<https://www.dropticks.com/articles/architecture-guide-self-hosted-inference>

Related links

<https://docs.vllm.ai/>

<https://huggingface.co/docs/text-generation-inference/index>