

ML LABS INTELLIGENCE

How to Scale AI Architecture for Growth

AI systems that work at today's load often fail when usage, data volume, or workflow complexity grows. This article shows the architecture patterns that scale without brittle rewrites.

TECHNICAL

Author Omar Trejo

Date 2026-01-20

ML LABS

mllabs.com/intel/scaling-ai-architecture-for-growth

Your AI system works at current scale. But growth is coming, and the architecture that handles today's load won't handle tomorrow's. AI workloads are compute-intensive, data-hungry, and non-deterministic — inference bottlenecks, feature computation overhead, and cost curves that scale linearly with load will turn a manageable system into an operational crisis without architectural intervention. The patterns that work for web applications apply but aren't sufficient. AI at 10x requires decisions that anticipate failure modes specific to machine learning at scale.

Where AI Architectures Break

Three categories of failure dominate AI scaling.

Inference bottlenecks. Model serving that handles 100 requests per second collapses at 1,000. GPU memory becomes the constraint, batching strategies that worked at low volume create unacceptable latency at high volume, and cold-start times for model loading become visible to users.

Data pipeline failures. Feature stores, training pipelines, and data preprocessing that run fine with gigabytes fail with terabytes. Batch processing windows exceed available time. Data freshness requirements that were aspirational become critical.

Cost explosions. AI infrastructure costs that scale linearly with load produce manageable budgets at current scale and terrifying ones at 10x. Without architectural changes, a \$50K/month AI infrastructure bill becomes \$500K/month — often without proportional revenue growth to justify it.



The non-ML components of an ML system – data pipelines, serving infrastructure, monitoring, and configuration – account for the majority of code and scaling bottlenecks. The model itself is rarely the problem.

Research on hidden technical debt in ML systems (NeurIPS, 2015) documented this finding across Google's production fleet, and **a survey of ML deployment challenges** (ACM Computing Surveys, 2022) confirmed it remains the dominant pattern across industries.

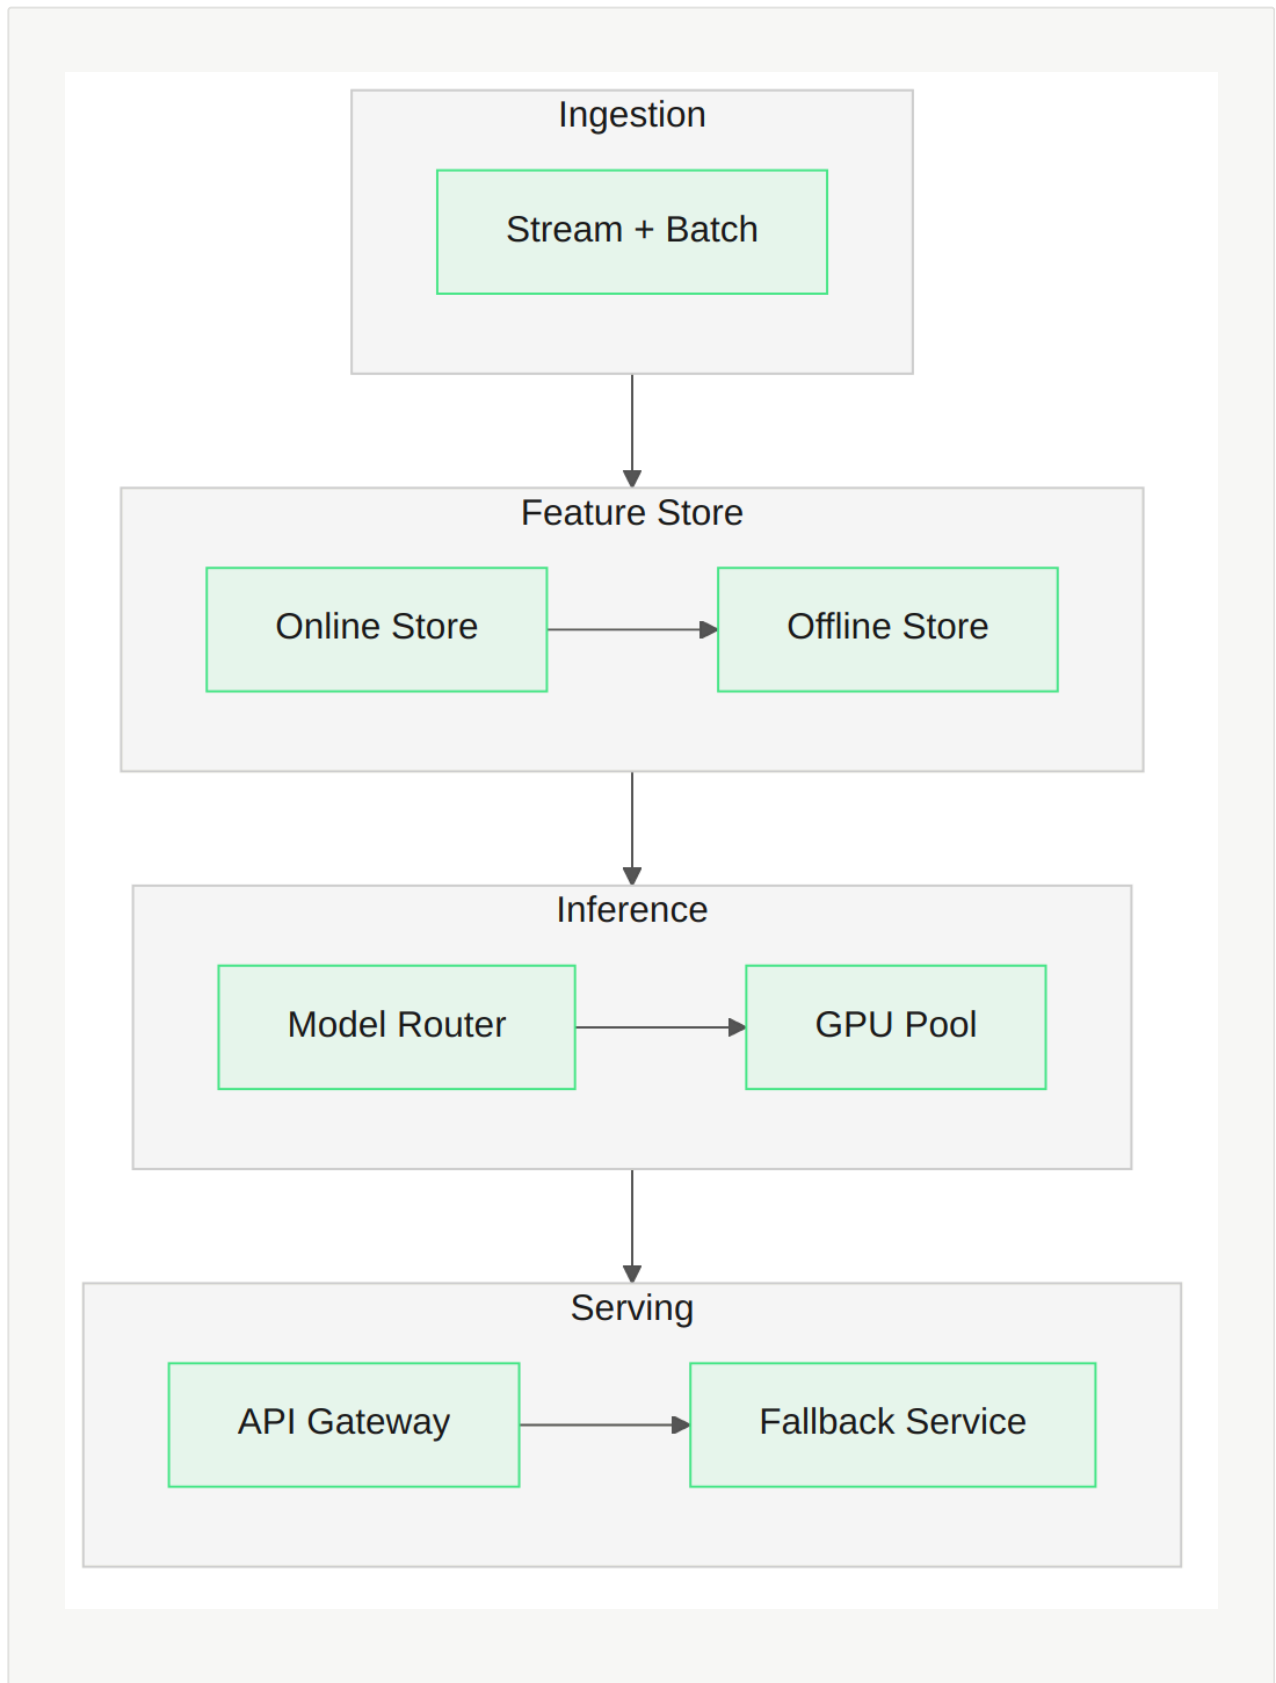
The 10x Stress Test

Run this diagnostic across four dimensions before designing for scale:

- **Inference latency.** What happens to P99 latency when you 10x the request rate? Degradation here hits users first.
- **Pipeline throughput.** Can feature computation handle 10x data volume in the same processing window? Batch jobs stall first.
- **Cost projection.** Multiply current infrastructure costs by 10 — is that sustainable, or does the unit economics break?
- **Failure blast radius.** If one component fails at 10x load, does it cascade or does the rest of the system stay healthy?

If batch processing already takes 20 hours, 10x volume means it won't complete within the same window. These diagnostics tell you where to invest first.

Scalable AI Architecture Patterns



Pattern 1: Decoupled Feature Computation

The most common bottleneck is computing features at inference time. If your model needs 50 features and each requires a database lookup, latency scales linearly with feature count.

The solution is a **dual-layer feature store**: an online store (Redis, DynamoDB) for single-digit millisecond reads, and an offline store (Parquet on S3, BigQuery) for training and batch computation. Features are precomputed and materialized so inference requires only lookups. Critically, both stores must use the same computation code — [Uber's Michelangelo platform](#) (Del Balso & Hermann, 2017) demonstrated that training-serving skew is a silent killer that must be eliminated architecturally. [Research on data lifecycle challenges in production ML](#) (ACM SIGMOD, 2018) formalized this pattern.

Pattern 2: Intelligent Model Routing

Not every request needs your most expensive model. A model router directs requests to the appropriate model based on complexity, latency, and cost constraints.

Simple requests route to small CPU-based models; standard requests to shared GPU instances; complex requests to larger dedicated models. The router itself can be a lightweight ML model trained on request characteristics. [Research on intelligent model cascading](#) (Chen et al., 2023) showed this reduces compute costs by up to 98% while maintaining quality.

Pattern 3: Stateful Horizontal Scaling

AI inference is harder to horizontally scale than stateless services because models carry state. Three strategies address this: model sharding across multiple GPUs, replica-based scaling behind a load balancer with health checks, and serverless

GPU inference for variable-load workloads. [Research on memory-efficient LLM serving](#) (SOSP, 2023) demonstrated that PagedAttention achieves 2-4x throughput improvements over prior approaches.

Pattern 4: Graceful Degradation

At 10x scale, failures are routine. Three mechanisms provide resilience: (1) circuit breakers that switch to a fallback when a model endpoint fails or exceeds SLAs, (2) priority queues so revenue-critical requests are served before background tasks, and (3) load shedding that deliberately drops low-priority requests rather than degrading all traffic. [Load shedding reference architecture](#) (AWS, 2024) demonstrates that proactive shedding maintains P99 latency even under 3x expected load.

Pattern 5: Cost Management at Scale

Three techniques produce sub-linear cost scaling: (1) spot instances for checkpointable training workloads, (2) mixed-precision inference in FP16/INT8 which reduces GPU memory 2-4x ([TensorRT benchmarks](#) (NVIDIA, 2024) demonstrate minimal accuracy loss), and (3) semantic caching to eliminate redundant inference requests. Multi-tenancy via GPU time-slicing across multiple models further improves utilization.

When This Approach Does Not Apply

These patterns assume operational maturity. Without comprehensive observability — latency distributions, error rates, resource utilization per model — adding architectural complexity creates more problems than it solves. If a model deployment failure today requires significant manual effort to diagnose and roll back, invest in deployment automation before scaling architecture.

First Steps

1. **Load test at 3x, 5x, and 10x.** The first component that breaks is where architectural intervention delivers the highest ROI.
2. **Implement a feature store** if you're computing features at inference time — almost always the highest-return scaling investment.
3. **Add graceful degradation.** Define fallback behavior for every model endpoint and test it under load.

Practical Solution Pattern

Decouple feature computation from inference, implement intelligent model routing, and define explicit degradation behavior for every model endpoint before load spikes occur. Start by load testing at 3x and 5x to identify the first bottleneck, then address it with the appropriate pattern before adding broader complexity.

This produces sub-linear cost scaling because each pattern targets a specific failure mode. Precomputed features eliminate per-request overhead; routing ensures expensive models handle only requests that require them; degradation keeps high-priority traffic within SLA when capacity is saturated. If the system is already live and the next risk is architectural, an **AI Technical Assessment** can map scaling constraints before they turn into outages or rewrites.


References

1. Sculley, D., et al. **Hidden Technical Debt in Machine Learning Systems.** *NeurIPS*, 2015.

2. Paleyes, Andrei, Raoul-Gabriel Urma, and Neil D. Lawrence. **Challenges in Deploying Machine Learning: A Survey of Case Studies**. *ACM Computing Surveys*, 2022.
3. Polyzotis, Neoklis, et al. **Data Lifecycle Challenges in Production Machine Learning**. *ACM SIGMOD Record*, 2018.
4. Chen, Lingjiao, et al. **FrugalGPT: How to Use Large Language Models While Reducing Cost and Improving Performance**. *arXiv*, 2023.
5. Kwon, Woosuk, et al. **Efficient Memory Management for Large Language Model Serving with PagedAttention**. *SOSP*, 2023.
6. Amazon Web Services. **Using Load Shedding to Avoid Overload**. *AWS Builders Library*, 2024.
7. Uber Engineering. **Meet Michelangelo: Uber's Machine Learning Platform**. *Uber Engineering Blog*, 2017.
8. NVIDIA. **TensorRT**. *NVIDIA Developer*, 2024.

PATTERNS

LOAD DISTRIBUTION DIAGRAMS



AMS

LOAD PATTERNS



ML LABS

Custom AI Systems for High-Value Workflows

mllabs.com