

## Architectural Patterns for Integrating Large Language Models (LLMs) into Node.js Server Applications

Oleksandr Tserkovnyi  
TrialBase Inc., Principal Engineer  
Dominican Republic, Punta Cana

### ABSTRACT

This article examines the evolution and systematization of architectural patterns for integrating large language models (LLMs) into server applications built on the Node.js platform, against the backdrop of the rapid diffusion of generative technologies in industrial software development and the expanding market for Retrieval-Augmented Generation (RAG) solutions. The relevance stems from the fact that by 2025, LLMs will have become an indispensable component of digital products, while server architectures must embed computational speech into existing infrastructures under constraints of token budgets, call costs, and network latency. The objective is to identify and analytically describe stable architectural patterns that enable efficient, predictable LLM integration in Node.js backends. Methodologically, the work combines systemic architectural analysis, modeling of interactions with LLM APIs, and content analysis of industrial practices, enabling the author to construct an engineering-economic efficiency model for each configuration. The article's novelty lies in formulating the concept of a balanced LLM-integration architecture in which throughput, token price, and service-layer observability are treated as interdependent architectural variables. An evolutionary pathway is proposed for transitioning from monolithic model calls to microservice and serverless patterns, informed by market growth dynamics and the scaling of compute resources. The article will benefit researchers and engineers engaged in server-application architectural design, cloud-service developers, and AI-engineering specialists aiming for resilient and cost-balanced deployment of LLM technologies in production environments.

**Keywords:** large language models, Node.js, architectural patterns, microservices, serverless computing

### INTRODUCTION

By 2025, large language models will have effectively become standard in industrial development. According to the annual Stack Overflow survey, over 80% of practicing developers have already used the OpenAI GPT family in work projects, with this share nearly doubling over the past two years (Stack Overflow, 2025). Such high penetration gives rise to a new engineering expectation: the server must be able to converse, which means that any modern Node.js backend inevitably faces the task of embedding LLM functionality as naturally as databases or message queues once did.

The reason is not mere trend-following; it is product economics. The market for retrieval-augmented generation (RAG) systems is already estimated to be nearly \$2 billion, and according to Mordor Intelligence, it is growing at a compound annual rate exceeding 39% through 2030 (Mordor Intelligence, 2025). For businesses, this implies that the costs of implementing an LLM layer are increasingly recouped not by quarter's end but by sprint completion: models reduce support time, improve support conversion, and create new paid knowledge-synthesis services. Consequently, for backend developers, early design of

architectural extension points that can accept external LLM APIs or local weights without requiring capital refactoring is advantageous.

In practice, needs crystallize into three task families. First, conversational interfaces, where an LLM generates replies or the user acts, are typically implemented in the Node.js ecosystem via streaming SSE to deliver text with negligible perceived delay. Second, final-document generation, reports, conversation summaries, and automated instructions, where asynchronous execution via a job queue is favored to avoid blocking HTTP traffic. Third, RAG services, which combine vector search with the model and expose an API for answers, are rapidly becoming a mandatory component of enterprise knowledge, as corroborated by the aforementioned market dynamics.

Architecture, however, encounters the physical limits of LLM APIs. First, token limits: even the mini version of GPT-4o supports a 128,000-token context window, which is appealing for documents yet requires strict input text normalization, or costs scale nonlinearly (OpenAI, 2024). Second, cost: the same GPT-4o mini is roughly sixty cents per million input tokens and two dollars forty cents per million output tokens. Hence, each superfluous uncached call translates directly into an invoice (OpenAI, 2025a). Finally, latency: even in an optimized stack, OpenAI states that GPT-4.1 mini halves latency relative to GPT-4o, yet the timescale remains seconds rather than milliseconds, compelling front-end designs that support progressive output delivery and timeouts for long contexts (OpenAI, 2025b). Thus, competent LLM integration in Node.js is a balance between linguistic power and engineering restraint, attainable only by accounting for real-world figures of tokens, dollars, and milliseconds.

## MATERIALS AND METHODOLOGY

### Research Methodology

The study of architectural patterns for integrating LLMs into Node.js server applications relies on a systematic analysis of contemporary engineering and economic sources, including model vendor documentation, industry reports on market dynamics, cloud provider practices, and empirical observations from open repositories. Eight primary sources were analyzed, spanning three thematic blocks: technological evolution of LLM APIs and their constraints (OpenAI, 2024; OpenAI, 2025a; OpenAI, 2025b), practices of quota management and compute scaling in clouds (Microsoft Learn, 2025), and trends in microservice architecture and streaming generation (The Business Research Company, 2025; LangChain, 2024).

The theoretical foundation comprises studies on RAG market growth, demonstrating a shift from experimental LLM integrations to robust product scenarios in which search and generation are unified within a single server boundary (Polan, 2025). According to Mordor Intelligence (2025), the RAG market is growing at an annual rate of over 39%, creating a compelling objective to devise architectural templates that optimize latency and call costs. As a practical framework, the work utilized OpenAI API performance and pricing metrics (OpenAI, 2025a), including input/output token prices, context limits (128,000 tokens), average generation latency, and worker load distribution. These parameters enabled the construction of an engineering-economic model, wherein each architectural pattern is evaluated by composite efficiency, which is the ratio of performance ( $\text{throughput} \times \text{latency}$ ) to operational expense (token cost  $\times$  invocation rate) (Aryan et al., 2023).

Methodologically, the study combines comparative architectural analysis, experimental modeling of Node.js service interactions with LLM APIs, and content analysis of industrial deployments. First, six baseline integration topologies were mapped: direct in-process model call, proxy API, job queue with workers, retrieval-augmented generation (RAG) layer, action-coordination (function-calling hub), and serverless edge deployment. For each topology,

characteristics were recorded: mean response time, degree of process isolation, sensitivity to traffic growth, and token volume per transaction.

The criteria for selecting and comparing patterns in this article are defined by strictly measurable engineering and economic variables, as well as the requirement for a balance between throughput and latency, token price, component isolation, and observability maturity. The final suitability of the configurations is assessed through composite efficiency as the ratio of throughput (throughput  $\times$  latency) to operational costs (token cost  $\times$  invocation rate).

The testing and performance recording parameters in this paper include average response time, degree of process isolation, sensitivity to traffic growth, and token volume per transaction, using OpenAI performance and pricing metrics (including a context limit of 128,000 tokens and generation latency) and cloud provider quota restrictions. The sample limitations stem directly from the analysis's reliance on primary sources and public quota policies, which do not provide detailed breakdowns by hardware configuration, Node.js version, and network RTT.

## Theoretical Definitions

Software architecture in contemporary theory is construed as the aggregate of a system's key structures, comprising elements, the relations among them, and the rules governing their composition; decisive, however, is not the mere inventory of modules, but those design decisions that constrain the space of subsequent change and thereby determine the controllability of product evolution (Wan et al., 2023). In this sense, architecture functions simultaneously as a descriptive model (what is connected and in what manner) and as a normative frame (what may be altered without compromising integrity). At the same time, architectural quality manifests as the system's capacity to preserve specified properties under increasing load, escalating scenario complexity, and shifting external dependencies. Such an interpretation binds architecture to the practice of trade-off analysis: target quality attributes, such as latency, throughput, reliability, security, and maintainability, are not added at the end, but compete with one another at the level of structural decisions; consequently, architectural choice is more appropriately understood as optimization under constraints rather than as a search for a singular correct solution. For this reason, applied architectural engineering assigns particular importance to the discipline of explicitly documenting decisions and their consequences, including scenario-based analysis of quality attributes and the managed evolution of architecture as a continuous process rather than a one-off design stage (Lytra et al., 2019).

The key concepts required for a systematic description of architectural decisions are architectural style and architectural pattern. An architectural style specifies a general mode of component decomposition and interaction (e.g., service decomposition, event-driven interaction, layered organisation), thereby defining typical boundaries and permissible dependencies (Esparza-Peidro et al., 2024). A pattern, in turn, fixes a stable configuration of structural decisions for a recurring problem, including characteristic trade-offs and conditions of applicability (Farshidi et al., 2020). It is essential to distinguish a pattern from an implementation technique: a pattern does not describe a specific library or protocol, but rather a persistent structure (for example, offloading computationally intensive operations into an asynchronous path) that can be realised by different means. Work on evolutionary architecture emphasises that under high requirements volatility, architectural decisions must be verifiable over time: mechanisms are needed that not only permit change, but also protect critical system properties from imperceptible degradation, thereby transforming architecture from a static blueprint into a managed system of constraints and feedback (Chondamrongkul & Sun, 2023).

Decomposition into autonomous services is typically described as an architectural style in which independent deployment and contract-based interface alignment are used to scale

development and reduce change risk; however, such independence is attained at the cost of more complex distributed interactions, the emergence of network uncertainty, and the need for strict observability (Azra, 2024). Academic surveys of microservices emphasize that gains in modularity and localization of change are accompanied by increased requirements for dependency management, resilience to partial failures, and request traceability across service boundaries; therefore, such systems require architectural coordination mechanisms, including operation idempotency, compensating actions, and explicit management of the state of long-lived processes (Söylemez et al., 2022). Observability in this context should be understood not as the presence of logs, but as the ability to reconstruct causal relationships among events, latencies, and failures in a distributed system from operational data. Studies of industrial practice in distributed tracing indicate that the very complexity of interpreting traces and aligning observability tooling becomes one of the constraints on effective microservice operations (Li et al., 2021).

Serverless computing constitutes a distinct deployment style in which the unit of execution is a function, while the platform assumes scaling, scheduling, and a substantial portion of operational responsibilities (Ghorbian & Ghobaei-Arani, 2025). In theory, this reduces barriers to entry and improves economics under variable load, but introduces specific constraints, most notably cold starts, upper bounds on execution time, and dependence on the runtime environment. Research works emphasize that the serverless approach is rational where it is beneficial to transfer resource management to the platform and pay for actual consumption; however, for long-running operations and complex context preparation, an architectural separation becomes necessary: heavy stages are moved into asynchronous paths, while boundary functions are responsible for rapid request intake, response formation, and correct enforcement of time limits (Toosi et al., 2024). Accordingly, serverlessness should be interpreted not as a universal replacement for servers, but as a means of shifting the boundary of responsibility, which necessitates rigorous design of the computation life cycle and explicit control of failures and retries.


The integration of large language models into server systems can be succinctly described through the concepts of computational dialogue, context, and constrained resource budgets. Context refers to the totality of input data provided to the model to generate an answer (Legashev et al., 2025). In practical operation, it is expressed in tokens and thus becomes a measurable resource that competes with latency and execution cost (Han et al., 2024). A consequential theoretical implication follows: the language-model layer within an architecture should be treated as an external computational subsystem with its own quotas and price, rather than as an ordinary library; therefore, architectural decisions must include control of context volume, reuse of results, management of call parallelism, and observability of costs at the transaction level. Retrieval-augmented generation is defined as a class of architectures in which the generative response is formed not only from the model's parameters, but also from retrieved knowledge fragments drawn from an external corpus (Klesel & Wittmann, 2025). In the basic scientific formulation, this is a combination of parametric memory and non-parametric memory in the form of an indexable store, which enables knowledge updates without retraining and links responses to sources, while simultaneously shifting part of the responsibility for result quality onto retrieval quality and the discipline of index management. As a result, an appropriate theoretical framework for such systems relies on architectural categories of boundaries, contracts, quality attributes, and managed evolution: the language model becomes one component whose behavior must be stabilized by engineering guardrails, while reproducibility, observability, and cost become as primary characteristics of architecture as latency and fault tolerance.

## RESULTS AND DISCUSSION

A sound LLM-service integration architecture begins with strictly measurable parameters; otherwise, discussion of templates devolves into guesswork. Foremost are throughput and latency (Dhaouadi et al., 2021). Even default access to GPT-4 on Azure yields 450,000 tokens per minute and 2,700 requests per minute, about 45 RPS at an average 100 tokens per request. These are default rate limits that Microsoft assigns to Azure subscriptions based on model and region. Microsoft's public documentation doesn't provide breakdowns like which GPU/CPU, which Node.js version, what RTT because these are service capacity allocation policies, not client SDK measurements. Moving to an enterprise tier increases limits by an order of magnitude but entails contractual procedures and delicate quota orchestration (Microsoft Learn, 2025). Therefore, if long documents must be served, the architecture must include streaming delivery and background job queues; for short chat sessions, a microservice with an SSE stream suffices.

Costs are measured not in virtual cores but in tokens. Hence, an economical scheme includes aggressive prompt-hash caching, elision of repeated history segments, and progressive reliance on cached input. The financial model directly dictates the technical one: any solution that fails to reduce total token flow is destined to fail the FinOps audit.


Isolating business logic from the AI core is no longer an architect's luxury, but a market necessity. The global microservices solutions market is expected to grow from USD 7.4 billion in 2025 to USD 15.6 billion by 2029, at a 20.6% CAGR, as shown in Figure 1 (The Business Research Company, 2025). According to the TBRC methodology, market assessments and forecasts are formed based on gold-standard sources and comparison/verification of market values through company financial indicators, statistical modeling (correlations, regressions, extrapolation, etc.) and mandatory expert validation (including interviews). The forecast is additionally adjusted taking into account macrofactors, including GDP, and qualitative assumptions.



**Figure 1. Global Microservices Solutions Market Forecast  
(The Business Research Company, 2025)**

Such tempo is impossible without independent release cycles: the LLM layer changes weekly, while domain services follow quarterly versions. A crisp boundary is not only the ability to patch without stopping the monolith but also insurance against regressions: rolling back one container is simpler than rebuilding an entire process.

Finally, verifiability and observability transform a diagram into a working product (Waseem et al., 2021). Field statistics indicate that 39.8% of teams prefer offline evaluation of model outputs, whereas only 32.5% use online metrics, real-time is harder to control than batch logs, as shown in Figure 2 (LangChain, 2024).



**Figure 2. Adoption of Control Mechanisms for AI Agents (LangChain, 2024)**

Estimates from TechSci Research (2024) show similar forecasts: The Global Microservices Architecture Market was valued at USD 5.67 billion in 2023 and is expected to reach USD 14.85 billion by 2029.

Hence, the requirements include tracing every call, collecting token metrics, and automatically performing A/B splits across model versions. Lacking these, the DevOps loop chokes, and incidents disappear into a black box that must be probed manually.

Ultimately, pattern selection seeks balance among four vectors: mechanical throughput, monetary token price, component isolation, and observability maturity. Ignoring any one causes the others to drift beyond tolerable bounds, pushing the project off schedule.

Direct in-process API call. Integrating the model's API within a single process is tempting for its minimal coupling. A controller receives the request, formats it in a service, and, without intermediate network hops, invokes the provider's client library. There is no serialization, no additional routing layers so that a feature can be implemented in just a few lines. The server maintains a familiar directory structure; tests are mocked using standard methods; deployment follows the same pipeline as the rest of the code.

The proximity has a flip side. Any model version update entails a full application release; vendor dependency hardens, making provider switching or even pricing plan changes more onerous than prototyping suggests. Load shares the same thread pool as web handlers. Moreover, security criteria are intertwined with domain logic because access tokens and billing counters are located within the process environment.

Nevertheless, this option is justified when the load is low and time-to-market is a primary concern. In a corporate prototype or early-stage startup, repackaging into a microservice often appears to be premature optimization; validating the hypothesis and estimating the per-request

cost is more important. If traffic volume is measured in tens of calls per minute and the model is only needed for brief outputs, the monolithic path remains the most rational.

Proxy service. Once external clients or multiple internal teams begin to consume the server, tight coupling to the model impedes growth. Extracting AI-communication logic into a dedicated edge layer turns it into an autonomous contract. This layer accepts external requests in a frontend-friendly format, most often a text interface over a representation transfer protocol, and immediately opens an output stream, allowing the user to view results as they are generated. The edge service becomes the central point for token accounting, limit enforcement, and prompt versioning; there is no need to rebuild other services, one container update suffices.

Deployment follows cloud-standard practice: build a container image, bake the dependency base into a layer, and ship it to an orchestrator cluster. Being small, the service initially requires a single replica; horizontal scaling can be achieved through an autoscaler. As load rises, new instances surface automatically; process isolation prevents mutual interference between text generation and core business functions.

The separate ingress also enables a strategic cache. Model responses, especially those grounded in static system prompts, often repeat. Storing results in fast memory reduces token spend and eliminates latency. Above the cache, a rate limiter protects against abuse that would turn the service into a costly mining rig via request storms. These measures are introduced locally without touching neighboring microservices: business teams release on their cadence, while the AI team experiments with more efficient models without risking the broader landscape.

Queue with background workers. Gradual load growth and expanding scenarios necessitate separating instantaneous user calls from heavyweight operations on long documents. In practice, this leads to an intermediate job queue, where the edge layer places a task description. An asynchronous worker then consumes the message and calls the language model. This three-segment craft, ingress, message broker, worker, avoids holding connections open until an answer is ready and frees web threads for new clients. The larger the input, the greater the gain: long generation or post-processing becomes background work; the user receives a quick confirmation identifier and can later poll task status.

The message broker is pivotal. Choices include lightweight BullMQ, optimized for Node.js simplicity; classic RabbitMQ with flexible routing; and distributed Kafka, designed for end-to-end data streams. If the bet is thousands of parallel jobs and guaranteed delivery, a broker with a disk log and acknowledgments prevails. If minimal latency and simple deployment are the primary requirements, a memory-oriented solution built on top of a Redis cluster is sufficient. Regardless of brand, resilience requires durable queues and worker-restart policies in the event of abnormal termination.

As the project grows, workers scale horizontally. Rather than rebuilding images, simply bring up additional instances with the same code; the broker distributes messages. Note, however, that concurrent worker and generation-intensity increases sharply raise token expenses, so autoscaling must consider price metrics, not just CPU load. Rotation also touches model updates: roll out a new prompt or parameters to a fraction of workers to observe answer quality before a complete switchover, with the option to roll back changes without impacting other services.

RAG layer. When the application must answer questions grounded in private knowledge corpora, queues alone are insufficient. A retrieval-augmented layer emerges, in which the input text is first embedded, the nearest items are retrieved from a vector store, and the retrieved context is then fed to the model. The pipeline, encompassing ingest, indexing, retrieval, and generation, is designed to allow each stage to evolve independently. From a Node.js standpoint, a thin adapter over a generic dense-vector client suffices. Storage options include cloud-hosted

Chroma, fully managed Weaviate, or scalable Pinecone; the balance between storage cost and reverse-search speed primarily dictates the choice.

The index lifecycle, unlike a classic database, requires special care. As source documents update, old embeddings lose relevance; therefore, incremental refresh and versioning are necessary to maintain reproducibility of the generation. A common tactic is to maintain two indices: the active one serves queries while a reserve is rebuilt on new data, followed by an atomic role swap. This eliminates inconsistency windows and enables quality control by comparing model answers on old vs. fresh knowledge before production rollout.

Action coordination (function-calling hub). Shifting from passive generation to scenarios where the model orchestrates external actions requires a coordinator that binds language-engine intent to specific internal or third-party services. The coordinator accepts the model's response as a strictly defined JSON object, in which each function is declared by a signature, namely, its name, arguments, and expected result. Developers define these signatures once; thereafter, the Node.js node automatically matches the model-provided structure to implementations and invokes the required method. Eliminating manual parsing reduces transformation errors and accelerates the addition of new actions, as a new service requires only a declarative entry.


Complexity increases when the model requires multiple calls within a single session. The coordinator must persist intermediate states, check idempotency, and, where necessary, fold the chain into an atomic transaction. For example, if one must create an invoice and then confirm payment, the two steps form a logical pair: the failure of the second step must compensate for the first. The node implements this via an action log and compensation mechanisms akin to the saga pattern; even if the process crashes, the sequence can be recovered and completed correctly.

Practical value emerges in integrations. CRM interaction becomes a natural dialogue: the model requests card creation, the coordinator executes it, the new record is returned, and the model's next move composes a user notification. Payment is analogous: the model initiates an operation, receives a transaction ID, and either confirms the charge or reports failure. The user experiences a seamless conversation, unaware of numerous hidden network exchanges.

Serverless edge. When serving a global audience, network transit to a central data center can exceed the time required for text generation. At this point, a serverless approach, where the model-calling function is deployed as close to the client as possible, is preferable. Amazon Lambda, Vercel Edge Functions, and Cloudflare Workers run code at distributed network nodes, charging strictly for actual resource usage. This scheme immediately addresses geographic distance but introduces constraints: cold starts increase first-request latency, and maximum runtime is capped.

To minimize cold-start impact, functions are kept warm by small synthetic event trickles or by lazily loading dependencies after the first output tokens, creating the illusion of immediate response. For long messages, streaming results to the client stretches generation time without blocking the interface. Runtime limits are mitigated by process partitioning: heavy context construction executes in a background queue, while the serverless function formats the reply from prepared data.

Proximity to the edge enables a cache shared by instances within a region. Storing request-response pairs in a distributed key store reduces token costs and variance in answers on repeated queries (Gilbert & Lynch, 2002). Since the key derives from a normalized-prompt hash, even slightly different requests can resolve to a shared result after canonicalization. The developer selects an expiration strategy and an affordable memory budget; the platform scales the rest. This triad, edge function, cache, and queue, serves a global audience at near-network latency without sacrificing generative richness. The above patterns are systematized in Figure 3.



**Figure 3. Scaling AI Model Integration (compiled by the author based on his own recommendations)**

In applied practice, recommendations for integration cannot be rigidly tied to the Node.js + OpenAI pair, because competing providers and model delivery channels impose different constraints on protocols, inference parameters, observability, and risk management. For example, Anthropic's messaging interface explicitly provides a streaming mode via SSE, and stream implementations are available in standard toolkits for multiple languages, making it possible to reproduce the same user-facing properties (low perceived latency, incremental text output) outside the Node.js ecosystem as well (Claude Docs, n.d.). In addition, a variant occurs in which the model provider is not addressed directly; instead, the same model is consumed through a cloud marketplace of managed models: thus, Claude is available through Amazon Bedrock with formalized inference parameters, which changes the architectural framing of the task (Amazon Bedrock, n.d.). A unified point of quota enforcement, logging, and access control emerges at the level of the cloud platform. Still, dependence on a specific delivery environment and its constraints regarding regions, logs, and network paths is strengthened.

An alternative class of competing approaches is associated not so much with the provider as with the choice of execution platform and the degree of autonomy: some teams move the intellectual layer to Python/FastAPI to simplify the implementation of asynchronous scenarios, background operations, and integration with data-processing tools; in doing so, the standard background-task mechanism in FastAPI supports moving post-processing out of the synchronous response path, thereby bringing the architecture closer to the queue + workers pattern even without the immediate introduction of a separate broker (FastAPI, n.d.). Where serverless delivery is required, such a service can be packaged into a container image and

deployed in AWS Lambda, enabling lightweight regional request-ingress points and scaling according to actual load. However, the constraints of the runtime environment and cold starts must be taken into account (AWS, n.d.). Finally, as an alternative to fully managed models, the Llama family is increasingly considered, where self-deployment and control over data and latency are permitted; in this case, the architecture is typically separated into an inference layer (a dedicated server/endpoint for the model) and an application orchestration layer (for example, the same FastAPI or another gateway), and the choice becomes a compromise between predictability of per-call costs and capital expenditures for compute resources and operations (GitHub, n.d.).

Serverless becomes more expensive than containers when the LLM layer ceases to be a rarely called function and becomes a constantly loaded service: a steady stream of requests throughout the workday or around the clock, a high proportion of long streaming sessions, and regular background tasks that keep the execution busy for a significant amount of time. Under these conditions, the per execution cost begins to scale linearly with the number of calls and the duration of the execution, and you additionally pay for the inevitable overhead of starting and initializing the environment, loading dependencies, establishing connections, and recreating local caches, because function instances are short-lived and frequently recreated. The container model becomes economically preferable when you can maintain a small pool of constantly running instances, thereby reusing warmed-up dependencies, stable connections, local cache, and a controlled concurrency pool. It also becomes more cost-effective when you need predictable latency without the need for special always-on modes of serverless computing. If you are forced to maintain warm-up or reserve capacity to prevent cold starts, serverless effectively loses the advantage of variable pricing and approaches the cost of always-on capacity, while remaining more stringent in terms of execution time limits and state management. Finally, serverless often loses out in terms of total cost of ownership when the logic involves toolchains and complex orchestration with intermediate state: the need for external state storage, deduplication, idempotency, and retries increases the share of infrastructure operations per request, while containers allow you to implement the exact mechanisms with fewer external calls and denser process-level observability.

Infrastructure hosting the machine-intelligence layer is as fragile as prompt logic, so operational and financial questions must be addressed concurrently. In practice, provisioning code emerges first; modular resource descriptions capture cluster composition, worker image versions, and network parameters in a human-readable template. A single manifest change yields a predictable plan; a CI/CD pipeline attached to the repository wraps the plan into an artifact, signs it, and rolls it across environments. Binding a provisioning state store to immutable container tags yields determinism, as the same template consistently yields the same result in both test and production environments. Secret channels deliver model keys and broker tokens without requiring entry into the commit history, thereby minimizing leakage risk.

As applications begin to consume tokens actively, the scene shifts: accounting and cost forecasting take center stage. Compute providers throttle not only technically but fiscally, imposing per-minute character limits and monthly caps. To avoid blocks or uncontrolled spending, monitoring is configured with a dual-track approach: metric streams compute real-time consumption, and a scheduler extrapolates to period-end. If a forecast exceeds a predefined threshold, the system signals an operations chat or temporarily throttles call frequency. Such auto-regulation is less expensive than emergency credit line extensions.

The collection and storage of model-access logs adhere to data protection regulations. User text is trimmed to context-relevant fragments, and identifiers are anonymized before being stored in an indexable format. Encrypted backups, delayed deletion, and strict role-based access policies render logs an analytic asset without becoming a leakage point. During external audits,

it is sufficient to present rotation configurations and proof that personal fields do not extend beyond the contractually specified region.

Architecture evolves through several stages. A project begins as a monolith with direct model-library calls, shortest path, and thus, it is natural for idea validation. As parallel sessions grow, AI-communication logic moves to a separate proxy, shortening core-code release times and increasing security. Next come the queue and background workers: heavy jobs migrate from the synchronous path to asynchronous execution, accompanied by a RAG layer that injects private knowledge into the context. When the audience spans continents, the engine transitions to serverless edge functions: latency falls to the network-stack bound, and payment becomes strictly proportional to actual usage. Throughout, central orchestration, compensation logs, and budget sentinels persist, ensuring the new topology remains as governable as the very first line of code. Thus, LLM-integration resilience is determined not only by engineering patterns but also by fine-tuned token economics and data-handling modes, enabling scaling without budget shocks or regulatory collisions.

## CONCLUSION

Architectural patterns for integrating LLMs into Node.js server applications represent not merely another technological iteration but a shift to a new paradigm of server-logic design, where computational speech becomes as infrastructural as a database or message broker. Empirical findings and market metrics analyzed herein indicate that sustainable LLM deployment necessitates the concurrent management of three variables: throughput, token cost, and service-layer observability. Only in their concerted balance does a model transform from smart add-on into a structural architectural element, delivering measurable efficiency gains without degrading operational characteristics.

The analysis suggests that architectural dynamics follow a law of gradual complication, progressing from monolithic direct library calls to an isolated proxy, then to job queues, and ultimately to a distributed serverless infrastructure at the edge. Each turn of this evolution is justified not by aesthetic design preferences but by concrete metrics, rising parallel sessions, increasing context lengths, and geographic audience expansion. Simultaneously, the shift toward microservice modularity and asynchronous patterns becomes not only a technical but also an economic imperative: token price and response latency act as coequal variables in the architectural equation.

Therefore, integrating LLMs into the Node.js milieu cannot be reduced to a simple API call. It is a process of systemic coupling among engineering, the FinOps discipline, and data ethics, wherein every new feature inevitably becomes part of a governed ecosystem, with quotas, logs, version control, and privacy policy. The ultimate resilience of such architecture is determined not by peak performance but by its capacity to remain deterministic amid model updates, worker-pool growth, and provider changes. In this sense, LLM-integration patterns constitute a modern analogue of classic principles of dependable systems: component independence, infrastructure reproducibility, and cost predictability.

The practical result of this article is a systematization of robust integration topologies and a logic for selecting them as load and complexity increase. Direct invocation of a model within an application is only justified under low-traffic conditions and prioritizing production speed. Still, it strengthens coupling and blurs the distinction between update and security risks, as well as domain logic. Moving LLM communication to a proxy service creates a contract boundary. It centralizes token accounting, limiting, prompt versioning, streaming, and caching, allowing the LLM layer to be changed independently of the rest of the system.

For long documents and complex tasks, the author emphasizes the need for queues and background workers, which remove HTTP traffic blocks and enable horizontal scaling of processing. However, this approach requires consideration of token costs during parallelism

and autoscaling. For responses based on private corpora, a RAG layer is introduced, featuring separate indexing, retrieval, and generation stages, as well as an index lifecycle discipline, because source updates directly impact reproducibility and quality. For action scenarios, a function call coordinator with strict JSON contracts is described, which requires idempotency and compensation in multi-step chains.

For a global audience, a serverless edge is proposed as a means of reducing network latency; however, its limitations are emphasized, including cold starts and execution time limits, which necessitate the division of computations into a fast edge layer and a slower background stage. A key conclusion of the article is that operational maturity is integral to the architecture: call tracing, token metrics, model version control, and experimentation mechanisms are necessary, as well as reproducible infrastructure deployment, financial consumption forecasting, and data protection rules in logs. The author's final position is that LLM cannot be treated as just another library: it is an external computing layer with quotas and pricing, and resilience is achieved only through the simultaneous management of performance, costs, and observability.

Consequently, a mature approach to incorporating language models into server applications treats them not as an external service, but as an autonomous computational intelligence layer subject to the same laws of scaling, orchestration, and budget planning as other elements of an industrial backend. In this synthesis of engineering rigor and adaptive cognition lies the strategic significance of the architectural patterns considered for the next generation of distributed systems.

## REFERENCES

Amazon Bedrock. (n.d.). *Anthropic Claude Messages API*. Amazon Bedrock. Retrieved December 7, 2025, from <https://docs.aws.amazon.com/bedrock/latest/userguide/model-parameters-anthropic-claude-messages.html>

Aryan, A., Nain, A. K., McMahon, A., Meyer, L. A., & Sahota, H. S. (2023). The Costly Dilemma: Generalization, Evaluation and Cost-Optimal Deployment of Large Language Models. *Arxiv*. <https://doi.org/10.48550/arxiv.2308.08061>

AWS. (n.d.). *Deploy Python Lambda functions with container images - AWS Lambda*. AWS. Retrieved December 8, 2025, from <https://docs.aws.amazon.com/lambda/latest/dg/python-image.html>

Azra, J. M. A. (2024). Exploring Observability Design Patterns of Microservices: Challenges and Solutions. *International Journal for Multidisciplinary Research*, 6(2). <https://doi.org/10.36948/ijfmr.2024.v06i02.21600>

Chondamrongkul, N., & Sun, J. (2023). Software evolutionary architecture: Automated planning for functional changes. *Science of Computer Programming*, 230, 102978. <https://doi.org/10.1016/j.scico.2023.102978>

Claude Docs. (n.d.). *Streaming Messages*. Claude Docs. Retrieved December 5, 2025, from <https://platform.claude.com/docs/en/build-with-claude/streaming>

Dhaouadi, M., Spencer, K. M. B., Varnum, M. H., Grubb, A. M., & Famelis, M. (2021). Towards a Generic Method for Articulating Design-time Uncertainty. *The Journal of Object Technology*, 20(3). <https://doi.org/10.5381/jot.2021.20.3.a3>

Esparza-Peidro, J., Muñoz-Escoí, F. D., & Bernabéu-Aubán, J. M. (2024). Modeling microservice architectures. *The Journal of Systems and Software*, 213, 112041. <https://doi.org/10.1016/j.jss.2024.112041>

Farshidi, S., Jansen, S., & van der Werf, J. M. (2020). Capturing software architecture knowledge for pattern-driven design. *Journal of Systems and Software*, 169, 110714. <https://doi.org/10.1016/j.jss.2020.110714>

FastAPI. (n.d.). *Background Tasks*. FastAPI. Retrieved December 8, 2025, from <https://fastapi.tiangolo.com/tutorial/background-tasks/>

Ghorbani, M., & Ghobaei-Arani, M. (2025). Serverless Computing: Architecture, Concepts, and Applications. *ArXiv*. <https://doi.org/10.48550/arxiv.2501.09831>

Gilbert, S., & Lynch, N. (2002). Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. *ACM SIGACT News*, 33(2), 51. <https://doi.org/10.1145/564585.564601>

GitHub. (n.d.). *llama3*. GitHub. Retrieved December 9, 2025, from [https://github.com/meta-llama/llama3/blob/main/MODEL\\_CARD.md](https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md)

Han, T., Fang, C., Zhao, S., Ma, S., Chen, Z., & Wang, Z. (2024). Token-Budget-Aware LLM Reasoning. *ArXiv*. <https://doi.org/10.48550/arxiv.2412.18547>

Klesel, M., & Wittmann, H. F. (2025). Retrieval-Augmented Generation (RAG). *Business & Information Systems Engineering*, 67, 551–561. <https://doi.org/10.1007/s12599-025-00945-3>

LangChain. (2024). *LangChain State of AI Agents Report*. LangChain. <https://www.langchain.com/stateofaiagents>

Legashev, L., Shukhman, A., Badikov, V., & Kurynov, V. (2025). Using Large Language Models for Goal-Oriented Dialogue Systems. *Applied Sciences*, 15(9), 4687. <https://doi.org/10.3390/app15094687>

Li, B., Peng, X., Xiang, Q., Wang, H., Xie, T., Sun, J., & Liu, X. (2021). Enjoy your observability: an industrial survey of microservice tracing and analysis. *Empirical Software Engineering*, 27(1), 5507. <https://doi.org/10.1007/s10664-021-10063-9>

Lytra, I., Carrillo, C., Capilla, R., & Zdun, U. (2019). Quality attributes use in architecture design decision methods: research and practice. *Computing*, 102(2), 551–572. <https://doi.org/10.1007/s00607-019-00758-9>

Microsoft Learn. (2025, August 21). *Azure OpenAI in Azure AI Foundry Models Quotas and Limits*. Microsoft Learn. <https://learn.microsoft.com/en-us/azure/ai-foundry/openai/quotas-limits?tabs=REST>

Mordor Intelligence. (2025). *Retrieval Augmented Generation Market Size, Share & 2030 Growth Trends Report*. Mordor Intelligence. <https://www.mordorintelligence.com/industry-reports/retrieval-augmented-generation-market>

OpenAI. (2024, July 18). *GPT-4o mini: advancing cost-efficient intelligence*. OpenAI. <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>

OpenAI. (2025a). *API Pricing*. OpenAI. <https://openai.com/api/pricing/>

OpenAI. (2025b). *Introducing GPT-4.1 in the API*. OpenAI. <https://openai.com/index/gpt-4-1/>

Polan, S. (2025). Retrieval-Augmented Generation: Architecture, Techniques, and Evaluations. *Journal of Modern Technology and Engineering*, 10(1), 42–56. <https://doi.org/10.62476/jmte.10142>

Söylemez, M., Tekinerdogan, B., & Kolukisa Tarhan, A. (2022). Challenges and Solution Directions of Microservice Architectures: A Systematic Literature Review. *Applied Sciences*, 12(11), 5507. <https://doi.org/10.3390/app12115507>

Stackoverflow. (2025). *2025 Stack Overflow Developer Survey*. Stackoverflow. <https://survey.stackoverflow.co/2025/>

TechSci Research. (2024). *Microservices Architecture Market By Size, Share, and Forecast 2029*. TechSci Research. <https://www.techsciresearch.com/report/microservices-architecture-market/25049.html>

The Business Research Company. (2025). *Microservices Architecture Global Market Report 2025*. The Business Research Company.

<https://www.thebusinessresearchcompany.com/report/microservices-architecture-global-market-report>

Toosi, A. N., Javadi, B., Iosup, A., Smirni, E., & Dustdar, S. (2024). Serverless Computing for Next-generation Application Development. *Future Generation Computer Systems*, 164, 107573. <https://doi.org/10.1016/j.future.2024.107573>

Wan, Z., Zhang, Y., Xia, X., Yi, J., & Lo, D. (2023). Software Architecture in Practice: Challenges and Opportunities. *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. <https://doi.org/10.1145/3611643.3616367>

Waseem, M., Liang, P., Shahin, M., Di Salle, A., & Márquez, G. (2021). Design, monitoring, and testing of microservices systems: The practitioners' perspective. *Journal of Systems and Software*, 182, 111061. <https://doi.org/10.1016/j.jss.2021.111061>