

Enabling Efficient Serverless Inference Serving for LLM (Large Language Model) in the Cloud

Himel Ghosh*

School of Computation, Information and Technology, Technical University of Munich, Munich, Germany

Email address:

himel.ghosh@tum.de (Himel Ghosh)

To cite this article:

Himel Ghosh. (2026). Enabling Efficient Serverless Inference Serving for LLM (Large Language Model) in the Cloud. *Internet of Things and Cloud Computing*, 14(3), 56-67. <https://doi.org/10.11648/j.iotcc.20261403.11>

Received: 21 December 2025; **Accepted:** 4 January 2026 **Published:** 12 October 2026

Abstract: This review report examines the problem of cold-start latency in serverless inference for large language models (LLMs) and surveys recent systems that address it, with a particular focus on ServerlessLLM. In serverless environments, loading multi-hundred-gigabyte LLM checkpoints and initializing GPU resources can dominate end-to-end latency, making real-time applications such as conversational agents and translation services difficult to deploy cost-efficiently. The report first summarizes the cold-start problem and common mitigation strategies in today's Function-as-a-Service (FaaS) platforms, including container warming, host-side checkpoint caching, and cluster-level storage servers. It then reviews the design of ServerlessLLM, which exploits underutilized multi-tier storage in GPU servers, a loading-optimized checkpoint format, live inference migration based on token-level state transfer, and a startup-time-aware model scheduler. Additional methods from recent literature, such as RainbowCake's layered container caching, and the cold-start handling mechanisms of major cloud providers are also discussed. Drawing on the evaluation presented in the original work, the report highlights that multi-tier checkpoint loading and live migration reduce model startup time by up to $6-8\times$ compared to conventional formats, and deliver up to $28\times$ lower startup latency than Ray Serve for large models in realistic traces. The review concludes by outlining open challenges and future directions, including integrating model compression, multi-level inference pipelines, and advanced memory-offloading techniques to further improve the scalability and efficiency of serverless LLM inference.

Keywords: Serverless Inference, Cold-start Latency, Large Language Models (LLMs), Multi-tier Storage, Checkpoint Optimization, GPU Scheduling

1. Introduction

The advancement of serverless computing [32] in cloud technology has revolutionized how computational resources are accessed, enabling developers to focus on application functionality rather than infrastructure management[16]. In this paradigm, resources are dynamically provisioned by cloud providers, activating only when required and charging users based on utilization. However, serverless computing faces a significant hurdle when applied to Large Language Models (LLMs) [28], especially for applications like real-time assistants, translation tools, and chatbots that rely on swift response times. These models, due to their size-often reaching hundreds of gigabytes-and computational requirements, encounter delays due to what is known as

the cold-start problem [15]. This latency arises when serverless functions, previously idle, initiate, leading to delays from the loading of extensive Large Language Model (LLM) checkpoints and Graphics Processing Unit (GPU) resource activation. Such cold starts can significantly hinder performance in applications requiring real-time interaction, making solutions to this problem imperative for scalable, serverless LLM deployment.

This paper investigates the current trends in tackling the cold start latency in serverless inference and specifically focusses on ServerlessLLM by [13], an innovative system tailored to reduce cold-start latency. The proposed system leverages a multi-tier checkpoint loading mechanism that optimizes GPU memory usage, alongside a live inference migration protocol and an efficient model scheduler designed to minimize startup

time. ServerlessLLM enhances performance by making use of typically underutilized memory and storage resources within GPU servers, significantly reducing cold-start delays and improving cost efficiency.

This paper is structured as follows: Section 2 discusses the cold start issue in serverless LLM inference and presents an overview of current mitigation approaches. Section 3 introduces ServerlessLLM’s architecture and then reviews the methods originally introduced by [13]. The section also explores the necessity of live migration in LLM inference and the development of a migration strategy that ensures minimal interruption to model service. Section 4 presents the analysis of the evaluation results, demonstrating the efficacy of ServerlessLLM compared to traditional serverless frameworks. Section 5 discusses additional methods proposed by recent literature and discusses the Rainbowcake method [34] which employs a layered container architecture to reduce latency and resource wastage in serverless functions. By caching shared container layers, RainbowCake allows efficient, fine-grained resource management across functions, making it well-suited to high-concurrency environments. Section 6 offers an insight on how the Function-as-a-Service (FaaS) providers deal with cold starts. Finally, Section 7 offers potential future directions, including the integration of memory management optimizations and distributed resource allocation techniques, aiming to enhance serverless LLM scalability and efficiency further.




Figure 1. Serverless Computing Architecture [14].

1.1. Contributions

This paper provides a structured and comprehensive examination of cold-start latency in serverless inference for Large Language Models (LLMs), with emphasis on the architectural and systems-level mechanisms that determine model startup time. The main contributions are as follows:

- 1) Consolidated Survey of Cold-Start Mitigation for LLM Inference.

- 2) Analytical Review of ServerlessLLM Architecture.
- 3) Comparative Evaluation of Approaches Across Literature.
- 4) Identification of Open Problems and Future Research Directions.

1.2. Novelty and Distinction

Unlike existing surveys on serverless computing, which primarily address cold starts at the container or platform level, this paper focuses specifically on the unique challenges posed by serverless inference for Large Language Models (LLMs). We provide a mechanistic analysis of ServerlessLLM, examining multi-tier checkpoint loading, optimized storage formats, and token-state migration-and compare these techniques with both traditional FaaS optimizations and LLM-oriented systems such as FlexGen, ZeRO-Infinity, and RainbowCake. By unifying cold-start research from serverless systems, GPU scheduling, and large-model inference, this work offers a more integrated perspective than prior studies and highlights design trade-offs and open problems that existing reviews do not explicitly address.

2. Cold Start in Serverless Inference

Deployment of the LLMs on the serverless system often introduces significant latency due to cold-start issues. The cold start is caused by the time taken to initialize the GPUs in the serverless architecture with the LLM checkpoints, which take longer download times pertaining to their huge sizes in the hundreds of gigabytes [14, 15]. The LLM checkpoint with 130 GB (LLaMA-2-70B) from Amazon Simple Storage Service (S3) or blob storage has been reported to take at least 26 seconds using a network of 5GB/s. Then the loading of these checkpoints into the GPU followed by model initialization imposes further delay. Loading the LLaMA-2-70B onto 8 GPUs takes another 84 seconds. The tokens are generated within 100ms in the LLM inference process, this loading latency significantly slows down the process.[13]




Figure 2. Cold Start [27].

Cold Start Mitigation Approaches:

Some of the prevalent solutions to tackle the cold start issues are the following:

- 1) Over-subscription of the GPUs to cater to peak demand scenarios. Amazon Web Services (AWS) serverless inference keeps a number of GPUs in the warmed state to avoid cold starts. Although it is effective

for comparatively smaller models like BERT, it is not enough for LLM for their requirement of substantially high load of GPUs. [22, 31]

- 2) To avoid model downloads, model checkpoints are cached in the host memory of GPU servers. But this is effective till a few GBs and is completely inadequate for LLMs which exceeds hundreds of GB sizes. Host memory have a limited size leading to many cache misses in such cases resulting in frequent model downloads, thereby not resolving the issue. [11, 12, 17, 20]
- 3) Deployment of additional storage servers within the local cluster to cache model checkpoints, has been recommended. But this process eventually requires additional storage servers and high bandwidth network which significantly increase the costs thereby subverting the purpose of going serverless. [2, 15, 29, 35]
The use of network optimised AWS Elastic Cache servers to support a 70B model has been reported to a 100% cost increase. [13]

As is evident from the associated issues of LLM inferences in the traditional solutions, there is a strong need of new methods to tackle the cold start problem focussing the LLM requirements.

3. Other Methods

To address cold start delays in serverless computing, several methods have been outlined in [29], categorized broadly into approaches aimed at either minimizing the cold start duration (optimizing environments) or reducing the frequency of cold starts (pinging).

Container Preparation Optimization: [18, 19, 21, 26]

- 1) Persistent Containers: Platforms like OpenLambda and OpenWhisk pause containers after function execution to retain the container in memory, allowing rapid reuse if requests recur shortly after. However, this results in memory overhead.
- 2) Container Pools: Fission and Knative platforms maintain pools of always-ready containers, which can be immediately used for new requests, significantly reducing latency but at a higher resource cost.
- 3) Warmed Containers: Google Cloud Functions and AWS Lambda keep containers "warm" for a period after function execution, allowing for reuse within this timeframe if new requests arise. This approach helps reduce cold start frequency but has an associated resource wastage.
- 4) Hot and Cold Queues: Some platforms utilize dual-queue architectures, with "hot" containers ready for immediate use and "cold" containers as backup. Requests are initially handled by the warm containers, switching to cold only if needed, optimizing response time at the expense of resource consumption.
- 5) Application Sandboxing: The SAND platform groups functions of the same application within a single

container using application-level sandboxing. By isolating functions through processes rather than containers, container preparation delay is reduced, and shared libraries load only once per container, optimizing both time and memory. [1]

Library Loading Optimization:

Frequently used libraries are preloaded into memory to avoid repetitive loading delays. This technique, similar to SAND's approach, benefits applications where multiple functions share libraries, as they are loaded only once.

Pinging (Reducing Cold Start Frequency)

To reduce the frequency of cold starts, several systems implement periodic "pinging" of functions. By scheduling functions to run periodically using tools like CloudWatch, Dashbird.io, and Lambda Warmer, functions are kept "warm" by sending periodic triggers, thus preventing containers from becoming idle. [3, 9, 10]

RainbowCake, *rainbowcake* is an innovative cold-start mitigation method designed for serverless computing environments, focusing on reducing latency and resource wastage during function invocations. The method employs a layered container architecture, which allows for fine-grained caching and the sharing of container layers across different functions. Specifically, it introduces three distinct layers: Bare, Lang, and User. By maintaining idle containers and enabling partial keep-alive capabilities, RainbowCake allows different functions to share compatible layers, minimizing cold-start latencies associated with serverless functions. [34] offers significant reduction in cold-start latency through its layered caching strategy, which promotes resource efficiency by enabling container layer sharing among functions. Its lightweight design allows for scalability, effectively handling high concurrency and bursty workloads typical in serverless environments. However, the method's complexity in decision-making can hinder the optimization of trade-offs between latency and resource waste. Additionally, its reliance on accurate invocation predictions makes it sensitive to mispredictions, potentially leading to resource inefficiencies. Furthermore, while it aims to minimize operational overhead, managing multiple container layers can still introduce latency, especially in dynamic workloads.

The architecture includes two main components: the History Recorder, which captures invocation patterns using statistical modeling (e.g., Poisson distributions), and the Container Pool, responsible for managing the lifecycle of the containers based on predictive analytics of incoming requests. The dynamic scheduling of pre-warming and keep-alive strategies helps ensure that resources are allocated efficiently, addressing bursty workloads without incurring significant overhead.

4. FaaS Providers' Approach to Mitigate Cold Start

Amazon Web Services (AWS) provides Lambda as its primary FaaS offering, which is commonly used for various

serverless applications. For LLM inference, AWS Lambda is supported by additional services like SageMaker and Elastic Inference to address resource needs and cold starts. AWS introduced provisioned concurrency [24] to pre-warm a specific number of function instances, allowing them to be instantly available when a request comes in. This significantly reduces cold start times. AWS Elastic Inference allows users to attach just the right amount of GPU to an EC2 instance or SageMaker instance. This ensures that the GPU is ready when needed, reducing cold start times for GPU-bound tasks. AWS SageMaker enables serverless inference by automatically scaling compute resources based on traffic. It also uses model caching to avoid frequent reloading of large LLM checkpoints.

Google Cloud offers AI Platform and Vertex AI for managing machine learning workloads. Google Cloud Functions supports the use of pre-warmed instances [6–8] to mitigate cold starts. This is similar to AWS’s provisioned concurrency, where instances are kept ready to handle

incoming requests without initialization delays.

For LLM inference, *Microsoft Azure* integrates with services like Azure Machine Learning and Azure Kubernetes Service (AKS) to provide GPU-enabled environments. Azure Functions’ Premium Plan [4, 5] allows for pre-warmed instances that can be kept ready to handle requests instantly, thus reducing cold start times. Azure Machine Learning (ML) provides a scalable and optimized environment for running LLMs. It offers GPU-based inference and model caching to avoid frequent cold starts. Azure Functions supports custom Docker containers that allow users to configure environments ahead of time. Pre-loading large models in custom environments can reduce cold starts by ensuring that resources are initialized and ready when the function is triggered.

One common approach by each of these FaaS providers is to keep some GPUs already in the warmed state to tackle Cold start latency. This incurs additional costs and the task becomes more expensive.




Figure 3. Flowchart to illustrate the methods of the ServerlessLLM to mitigate Cold start problem.

5. Proposed Methods to Mitigate Cold Starts

A new method called ServerlessLLM has been proposed by [13] in their paper to address this particular issue. ServerlessLLM is a system designed specifically to reduce the cold start problem in serverless computing for Large Language Model (LLM) inference. It optimizes how models are stored, loaded, and managed, ensuring minimal startup latency.

5.1. Design Perspective

GPU servers are equipped with a multi-tier storage architecture offering substantial capacity and bandwidth. An 8-GPU server can handle up to 4 TB of main memory, 64 TB of Non-Volatile Memory Express Solid State Drives (NVMe SSDs), and 192 TB of SATA SSDs. However, in a serverless environment, a significant portion of the host memory and storage in these GPU servers often remains underutilized. Each GPU in these systems is linked to the host memory via a dedicated Peripheral Component Interconnect Express (PCIe) connection, ensuring a high combined bandwidth

between memory and GPUs. For instance, an 8-GPU server utilizing PCIe 5.0 can deliver a 512 GBps memory-to-GPU bandwidth, with approximately 60 GBps from NVMe SSDs to memory. The suggested approach leverages the unused multi-tier storage within the server to store models locally, improving load times and reducing latency. This method proves to be cost-effective due to the reuse of underutilized resources, scalable thanks to the system’s expandable storage and bandwidth, and sustainable, as future GPU servers will offer even greater capacities and bandwidth.

5.2. Fast Multi-tier Checkpoint Loading

ServerlessLLM introduces a multi-tier checkpoint loading system that accelerates large model loading by leveraging a hierarchical storage structure across GPU memory, Dynamic Random-Access Memory (DRAM), and Solid State Drives (SSDs). This approach minimizes loading bottlenecks and improves LLM initialization times.

- 1) Loading-Optimized Checkpoint Format: A new format facilitates sequential, chunk-based reading of model parameters, organized in partitions. This design reduces disk Input/Output (I/O) and accelerates data transfer.

Tensors (which represent the model's weights and parameters) are organized in partitions, allowing large chunks of data to be read sequentially from storage rather than one tensor at a time. This minimizes disk I/O operations and speeds up data transfers.

- 2) Multi-Tier Storage Hierarchy: GPU Memory serve as the final destination for inference. DRAM is the Intermediate cache for quick GPU access. SSDs are Long-term checkpoint storage for faster loading compared to remote sources.
- 3) Parallel Chunk-Based Loading: Data is loaded in parallel across storage tiers, maximizing bandwidth through multi-threading and overlapping operations.
- 4) Efficient Tensor Addressing: Each tensor's memory address is pre-computed, allowing for direct memory mapping. This eliminates unnecessary memory copy

operations and further speeds up the loading process.

5) Direct I/O and Pinned Memory:

The system uses Direct I/O (bypassing the operating system's cache) to read data directly into user space, ensuring more predictable and faster data access. Additionally, pinned memory is employed to allow direct data transfers between DRAM and GPU memory without involving the CPU, maximizing the efficiency of the PCIe bus and minimizing CPU overhead.

The system reduces cold start times significantly-up to 6X faster for small models (OPT-2.7B) and 8.2X faster for large models (LLaMA-2-70B) compared to PyTorch and SafeTensors. ServerlessLLM reaches near-maximum throughput from RAID0-NVMe SSDs, supporting predictable, high-speed data transfer ideal for real-time applications. [13]




Figure 4. The multitier storage design to mitigate cold start issues.

5.3. Review of Fast Multi-tier Checkpoint Loading

The method has many strengths, as following:

- 1) The reduction in loading times by 8.2X speed for large models is a substantial improvement. This is particularly impressive given the size and complexity of modern LLMs, which often exceed hundreds of gigabytes in size.
- 2) With the full use of the storage hierarchy and employing parallelism in data loading, the system maximizes resource utilization and reduces idle time. This makes serverless computing more resource-efficient for LLM workloads.
- 3) The method is scalable, as the system is able to saturate the bandwidth of storage devices, it can handle even larger models without introducing significant additional latency.

There are also some limitations to focus on and areas for improvement:

- 1) The set of assumptions made by the authors that the checkpoints have model execution files and model parameter files, may not always hold and in that case, the architecture needs to adapt to such specific scenarios or such approaches need to be standardized.
- 2) While the use of DRAM and SSDs as caching layers reduces the need for frequent accesses to slower remote storage, the system's performance is still constrained by the available memory resources. For extremely large models (e.g., those exceeding hundreds of GBs), there is a risk that the memory hierarchy could become a bottleneck, especially if multiple models are in use simultaneously.
- 3) Although the method is effective in reducing latency, the use of pinned memory, Direct I/O, and parallel I/O threads may increase the overall energy consumption of the system. This is a trade-off that needs to be carefully balanced, especially for applications where

power efficiency is a concern.

- 4) While the system is optimized for LLM inference, it may not generalize well to other types of AI workloads that have different memory access patterns or storage requirements. Additional work may be needed to adapt the system for tasks like real-time video processing or reinforcement learning, where the data access characteristics differ from those of LLMs.

5.4. Need for Live Migration in LLM Inference

Efficient LLM inference management is crucial when multiple models must run concurrently across servers with varied resource availability. The cold start issue, involving loading large model checkpoints into GPU memory, can significantly delay inference. We illustrate with a setup of two servers (Server 1 and Server 2) and two models (Model A and Model B), following [13].

Server conditions:

- 1) Server 1: Model A in DRAM, Model B in SSD, idle GPU.

- 2) Server 2: Model B in DRAM, Model A running on GPU.

Policies:

- 1) GPU Availability-Driven: Model B starts on Server 1 (available GPU, but Model B is in SSD), incurring high latency to load Model B into DRAM, though Model A remains unaffected.
- 2) Locality-Driven: Model B waits in a queue on Server 2, where it is already in DRAM, until Model A completes—delaying Model B’s startup.
- 3) Preemption-Driven: Model A is preempted on Server 2 for Model B, then reloaded on Server 1 from DRAM/SSD, leading to significant downtime for Model A.
- 4) Live-Migration-Supported Locality-Driven: Model A continues on Server 2 while preloading onto Server 1. Model B starts on Server 2, benefiting from DRAM locality, while Model A is seamlessly migrated to Server 1, minimizing latency for both models.

This live migration approach reduces latency, optimizes resource utilization, and scales efficiently with demand and model complexity.




Figure 5. Diagram that illustrates the Live Migration process.

5.5. Efficient Live Migration process

Several processes have been considered in [13] along with their demerits. Snapshotting LLM inference generates unnecessarily lengthy snapshots resulting in slower transfer times. Dirty-page based approach is not supported in GPU-enabled containers and VMs currently. [13] proposes to migrate minimal inference states to reduce network traffic. Instead of large Key-Value (KV)-Cache (1-10s GB), migration

of tokens (10-100s KB) is considered. Recomputing the KV-cache based on migrated tokens on the destination GPU is much faster. Also the destination server needs to rapidly synchronize with source server’s progress causing minimal migration times. For this, KV-cache for current tokens are recomputed on the destination GPU which is significantly faster than that in the source GPU. This has been implemented as a multi-round live migration process where the destination

server recomputes the KV-cache based on intermediate tokens sent by source server in each migration round and when the gap between inference states of source and destination minimises, the source stops generating and sends all tokens to the destination via request router and the further inference shifts to the destination server resulting in minimal interruption.

The process flow is as follows:

- 1) The scheduler sends a request to the dest server to load Model A into its GPU(s). However, if an idle instance of Model A already exists on the dest server, this step is skipped, and the process moves forward without loading.
- 2) Once Model A is successfully loaded, the scheduler sends a migration request to the src server. This request contains the address of the dest server, signaling that the migration process is about to begin.
- 3) Upon receiving the migration request, the src server marks itself as being in a "migrating" state. If the inference on Model A is incomplete, the src server sends a resume request to the dest server, which includes the intermediate tokens (the input tokens and the output tokens generated up to this point). If the inference is already completed, the src server simply notifies the scheduler without sending the resume request.
- 4) The dest server then processes the resume request by using the intermediate tokens to recompute the Key-Value (KV) cache. This step is crucial because it ensures that the dest server can seamlessly pick up the inference task from where the src server left off.
- 5) Once the resume request has been processed by the dest server, the src server halts its inference operations and sends a response back to the scheduler. This response includes all tokens generated during the migration process (i.e., the intermediate tokens and any additional tokens produced during the migration), along with a flag indicating that the task has been ?migrated.?
- 6) The scheduler finalizes the migration by unloading Model A from the src server and initiating the loading of Model B. Simultaneously, the request router checks the flag in the inference response. If the flag indicates that the task has been migrated, the router updates its routing table to replace the src server with the dest server. It then sends all tokens to the dest server, allowing the inference task to resume seamlessly on the new server.

5.6. Review of Live Migration

The autoregressive nature of LLM inference can lead to cases where a task completes on the source server (src server) before the migration fully completes. In such cases, the src

server finishes the inference task and informs the request router as usual. Additionally, it notifies the loading scheduler, which then instructs the destination server (dest server) to cease resuming and terminates the migration. If the src server fails during the loading phase (before step 2), the scheduler aborts the migration and unloads the model from the dest server. If the failure occurs during migration (steps 2-3), the scheduler directs the dest server to clear any resumed KV cache and unload the model. If the dest server fails during loading, the scheduler cancels the migration. If failure occurs while resuming, the src server notifies the scheduler of the failure and continues inference without interruption. By migrating inference tasks to servers with pre-loaded models, the process avoids the need for a complete model reload, effectively minimizing cold start times. This ensures that computational resources are more efficiently utilized, and models are able to resume inference quickly, even during migration. The ability to transfer only the token state and recompute the KV cache on the destination server allows inference to continue almost seamlessly. This ensures that the migration does not introduce significant interruptions or delays in processing, making it particularly useful for real-time LLM applications.

Although, the live migration process adds complexity to the system, particularly in terms of synchronizing between the source and destination servers. If inference completes during migration or there are timing delays, it may lead to inefficiencies or resource wastage, especially if unnecessary computations occur on the destination server. While recomputing the KV cache on the destination server reduces the need for transferring large amounts of data, it introduces computational overhead. This may still cause delays, especially for larger models, and can partially negate the latency savings from avoiding full checkpoint reloads. In case of server failures during migration, some resources may be wasted, particularly if significant amounts of computation have already occurred on the destination server before the failure is detected. Additionally, recovery from failures can introduce extra latency, affecting overall system performance.

5.7. Startup-Time-Optimized Model Scheduling

Another major cause of cold start latency is inefficient scheduling of inference tasks. In traditional systems, tasks are assigned to servers without considering whether the required model is already loaded. ServerlessLLM optimizes this process through startup-time-optimized model scheduling. The scheduler employs two components: model loading time estimator and model migration time estimator.




Figure 6. Estimation times for scheduling.

- 1) Model loading Time estimation: The model queuing time (q), model size (n) and the bandwidth (b) are considered. ServerlessLLM tracks bandwidth for network, SSD and DRAM and then model loading time is calculated as $q + n/b$
- 2) Model migration time estimation: For this, the model resuming time is calculated considering number of input prompt tokens (t_{in}), number of output response tokens (t_{out}), model specific parameters a and b that vary with each LLM batch sizes and other factors. Then the model resuming time is $a \times (t_{in} + t_{out}) + b$. However $t_{out} = d/t$ where d is the inference duration and t is the average time to produce a token, which are obtained by the scheduler querying the local request router.

ServerlessLLM evaluates all servers and selects the one that offers lowest estimated startup time. If no GPUs are free, the loading task is paused and retried once the request router informs the scheduler to release GPUs. [13] proposed a fail-safe process by utilizing a KV store to track server statuses. After server loading request, GPU status is loaded on the store, then after server confirms task confirmation, scheduler updates the storage status in the store and then notifying the request router which enables routing of the server. In failure event, latest server status from the KV store is utilized. The proposed scheduler is scalable as per current benchmarks. It is a fair process treating all models with equal priority and ensures that migrations don't impact latency.

Performance




Figure 7. Illustration of the performance of ServerlessLLM compared to baselines.

6. Evaluation and Results

The evaluation focuses on the effectiveness of loading-optimized checkpoints, overheads from live migration, and compares ServerlessLLM against large-scale serverless workloads based on real-world data. Two main testbeds are used: *GPU Server*: 8 NVIDIA A5000 GPUs, 1TB DDR4 memory, and 2 AMD EPYC Central Processing Units (CPUs), connected to a MinIO storage server. *GPU Cluster*: 4 servers connected via 10 Gbps Ethernet, each with 4 A40 GPUs, 512 GB memory, and PCIe 4.0 NVMe SSDs. Models like OPT, LLaMA-2, and Falcon in sizes (e.g., OPT-6.7B, OPT-13B, OPT-30B) are evaluated using datasets like GSM8K and ShareGPT for realistic input data. Workloads are derived from Azure Serverless Trace, reflecting real serverless inference conditions

6.1. Checkpoint Loading

ServerlessLLM achieves up to 8.2X faster model loading than PyTorch and Safetensors, notably reducing latency through chunk-based parallel loading and direct I/O. Unlike PyTorch, which copies data to host memory first, and Safetensors, which faces page faults, ServerlessLLM bypasses these overheads and consistently optimizes loading regardless of model type, focusing solely on checkpoint size. The system also loads LoRA adapters 4.4X faster and maximally utilizes storage bandwidth (e.g., RAID0-NVMe SSDs), significantly cutting cold start times in serverless environments. Bulk

reading, direct I/O, multi-threading, and pinned memory collectively improve throughput by mitigating small tensor reading delays, bypassing cache, and optimizing concurrency.

6.2. Model Scheduler

ServerlessLLM's cluster scheduler outperforms the default serverless scheduler and Shepherd* (a preemption-based variant), especially under high Requests Per Second (RPS) conditions with the OPT-6.7B model. ServerlessLLM's locality-aware, live migration approach reduces P99 latency by 1.95X over the default scheduler, which frequently reloads models from SSD. Compared to Shepherd*, ServerlessLLM achieves 2X lower P99 latency with ShareGPT by avoiding delays from checkpoint reloads in preemption. Even at peak GPU occupancy, ServerlessLLM maintains 35-45% lower P99 latency due to its efficient live migration, handling more migrations with minimal inference disruption.

6.3. Entire ServerlessLLM

In serverless deployment with Azure Serverless Trace, ServerlessLLM outperforms Ray Serve, Ray Serve with Cache, and KServe in distributed model serving. For OPT-6.7B, it initiates in 0.8 seconds, compared to 12.1 seconds (Ray Serve) and 8.2 seconds (Ray Serve with Cache). For OPT-30B, ServerlessLLM achieves a 7.5-second startup versus 213 seconds for Ray Serve, a 28X speedup. It maintains low latency across RPS ranges up to 1.1, achieving up to 212X

lower latency than Ray Serve variants. Resource-efficient, ServerlessLLM reaches 4-second latency with one GPU per server, while Ray Serve with Cache needs four GPUs for 12 seconds. Even optimized, KServe trails behind in Kubernetes cluster performance.

6.4. Remarks

Cold-start latency poses a major challenge in serverless systems, with various approaches being employed to mitigate the issue, such as fast image pulling, lightweight isolation techniques, snapshot and restore mechanisms, resource pre-provisioning, elastic resource provisioning, and process forking. However, these methods primarily target reducing container or Virtual Machine (VM) startup times without addressing the need to load large external states. Latest research has focused on optimizing cold-starts by accelerating model swaps between GPUs and host memory, but scaling these solutions for large language models (LLMs) remains a significant challenge. ServerlessLLM, in contrast, directly addresses the cold-start problem for LLM inference by introducing innovations such as optimized checkpoint formats and loading pipelines, live migration, and a cluster scheduler specifically designed for the demands of LLM workloads.

7. Future Scopes

The methods and evaluations presented by [13] demonstrate the system's effectiveness in minimizing cold start latency. While the current approach focuses on checkpoint optimization, integrating techniques like quantization and pruning—such as 4-bit compression as seen in FlexGen [25]—could further reduce memory usage and loading times for larger models. These optimizations could enhance the scalability of serverless systems with minimal impact on accuracy, especially for high-throughput tasks. Additionally, adaptive resource allocation mechanisms that dynamically adjust GPU, CPU, and memory based on real-time workloads could further minimize cold start delays. Auto-scaling with resource elasticity offers another promising avenue for improving responsiveness under varying load conditions.

Exploring cross-node offloading, where model states are distributed across multiple nodes, could also mitigate cold start latencies tied to single-node limitations, particularly for large LLMs that require more memory than a single GPU node can provide. Collaborative inference—distributing tasks across multiple devices or servers—could similarly reduce the model footprint per device, addressing cold starts. Systems like Petals [25] illustrate the potential of such approaches, though challenges remain in balancing latency and throughput across nodes.

FastServe [33] demonstrates the potential of preemptive scheduling at the token level to reduce job completion times (JCT) during LLM inference. Extending these strategies to optimize first-token latency while keeping overhead low could further enhance cold-start efficiency in serverless

LLM deployments. Current key-value cache management methods introduce overhead due to memory constraints, and innovations in memory-swapping techniques or distributed memory architectures could optimize resource utilization, reducing cold starts when memory is limited.

Systems like Tabi [30], which employ a multilevel inference engine that escalates from smaller models to LLMs only when necessary, offer a promising solution. Adapting this approach to serverless environments could preemptively handle simpler queries with lightweight models, reserving larger LLMs for more complex tasks, thus minimizing cold start latencies. Future research could focus on developing adaptive scheduling systems that automatically select the appropriate model based on query complexity, avoiding the use of heavy LLMs for tasks that require far less computation. Attention-based pruning techniques, as demonstrated by Tabi [30], could also be explored further to reduce computational overhead in serverless settings, enhancing responsiveness by minimizing model loading times and cold start delays.

ZeRO-Infinity's [23] memory-centric tiling and offload engine provide a compelling solution for efficiently utilizing heterogeneous memory systems, such as GPU, CPU, and NVMe. Applying similar strategies to serverless LLM inference could reduce cold start times by offloading less frequently accessed data to NVMe or CPU, easing GPU memory pressure. ZeRO-Infinity's dynamic prefetching and overlap-centric design [23], which overlaps data movement with computation, could also be implemented in serverless platforms to reduce wait times and ensure more seamless inference during high-demand scenarios. As serverless LLM workloads scale to trillions of parameters, techniques from ZeRO-Infinity-like NVMe offload and superlinear scalability—can help maintain performance while mitigating cold start issues.

ORCID

0009-0008-0492-6106 (Himel Ghosh)

Abbreviations

AI	Artificial Intelligence
AKS	Azure Kubernetes Service
API	Application Programming Interface
AWS	Amazon Web Services
BERT	Bidirectional Encoder Representations from Transformers
CPU	Central Processing Unit
DB	Database
DDR4	Double Data Rate 4
DRAM	Dynamic Random-Access Memory
EC2	Amazon Elastic Compute Cloud
FaaS	Function-as-a-Service
GSM8K	Grade School Math 8K
GPU	Graphics Processing Unit

I/O	Input/Output
JCT	Job Completion Time
KV	Key-Value
LLaMA	Large Language Model Meta AI
LLM	Large Language Model
LoRA	Low-Rank Adaptation
ML	Machine Learning
MLT	Model Loading Time
MMT	Model Migration Time
NVMe	Non-Volatile Memory Express
OPT	Open Pre-trained Transformer
P99	99th Percentile
PCIe	Peripheral Component Interconnect Express
RAID 0	Redundant Array of Independent Disks, Level 0
RPS	Requests Per Second
S3	Amazon Simple Storage Service
SATA	Serial Advanced Technology Attachment
SSD	Solid State Drive
VM	Virtual Machine
ZeRO	Zero Redundancy Optimizer

Author Contributions

Himel Ghosh: Conceptualization, Investigation, Writing – review & editing

Conflicts of Interest

The author declares no conflicts of interest.

References

- [1] I. E. Akkus et al., “SAND: Towards High-Performance Serverless Computing,” in *USENIX Annual Technical Conference*, pp. 923–935, 2018. Available: <https://www.usenix.org/conference/atc18/presentation/akkus> Accessed: Jan. 1, 2025.
- [2] Anyscale, “Loading LLaMA-2 70B 20× Faster With Anyscale Endpoints,” 2025. Available: <https://www.anyscale.com/blog/loading-llama-2-70b-20x-faster-with-anyscale-endpoints> Accessed: Jan. 1, 2025.
- [3] AWS, “Amazon CloudWatch,” 2025. Available: <https://aws.amazon.com/cloudwatch/> Accessed: Jan. 1, 2025.
- [4] Microsoft Azure, “Azure Tips and Tricks: Tip 260,” 2025. Available: <https://microsoft.github.io/AzureTipsAndTricks/blog/tip260.html> Accessed: Jan. 1, 2025.
- [5] Microsoft Azure, “Improving Azure Functions Cold Starts,” 2025. Available: <https://techcommunity.microsoft.com/t5/apps-on-azure-blog/our-latest-work-to-improve-azure-functions-cold-starts/ba-p/4164500> Accessed: Jan. 1, 2025.
- [6] Google Cloud, “Configuring Warmup Requests,” 2025. Available: <https://cloud.google.com/appengine/docs/legacy/standard/go111/configuring-warmup-requests> Accessed: Jan. 1, 2025.
- [7] Google Cloud, “Instance Autoscaling,” 2025. Available: <https://cloud.google.com/run/docs/about-instance-autoscaling> Accessed: Jan. 1, 2025.
- [8] Google Cloud, “Mitigating Cold Starts on Cloud Run,” 2025. Available: <https://medium.com/google-cloud/3-solutions-to-mitigate-the-cold-starts-on-cloud-run-8c60f0ae7894> Accessed: Jan. 1, 2025.
- [9] J. Daly, “Lambda Warmer,” 2025. Available: <https://github.com/jeremydaly/lambda-warmer> Accessed: Jan. 1, 2025.
- [10] Dashbird, “Dashbird Serverless Monitoring,” 2025. Available: <https://dashbird.io/> Accessed: Jan. 1, 2025.
- [11] D. Du et al., “Catalyzer: Sub-millisecond Startup for Serverless Computing with Initialization-less Booting,” in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '20)*, pp. 467–481, 2020, <https://doi.org/10.1145/3373376.3378512>
- [12] X. Liu et al., “FaaSLight: General Application-level Cold-start Latency Optimization for Function-as-a-Service in Serverless Computing,” *ACM Transactions on Software Engineering and Methodology*, 2023, <https://doi.org/10.1145/3585007>
- [13] Y. Fu et al., “ServerlessLLM: Low-Latency Serverless Inference for Large Language Models,” arXiv:2401.14351, 2024. Available: <https://arxiv.org/abs/2401.14351> Accessed: Jan. 1, 2025.
- [14] M. Ghorbian and M. Ghobaei-Arani, “A Survey on the Cold Start Latency Approaches in Serverless Computing: an Optimization-Based Perspective,” *Computing*, 2024, <https://doi.org/10.1007/s00607-024-01335-5>
- [15] M. Golec, G. K. Walia, M. Kumar, F. Cuadrado, S. S. Gill, and S. Uhlig, “Cold Start Latency in Serverless Computing: A Systematic Review, Taxonomy, and Future Directions,” *ACM Computing Surveys*, vol. 57, no. 3, pp. 1–36, 2024, <https://doi.org/10.1145/3700875>
- [16] IBM, “What Is Serverless Computing?” 2025. Available: <https://www.ibm.com/topics/serverless> Accessed: Jan. 1, 2025.

- [17] K. Kaffes et al., “Centralized Core-granular Scheduling for Serverless Functions,” in *Proceedings of the ACM Symposium on Cloud Computing*, pp. 158–164, 2019, <https://doi.org/10.1145/3357223.3362709>
- [18] P.-M. Lin and A. Glikson, “Mitigating Cold Starts in Serverless Platforms: A Pool-Based Approach,” arXiv:1903.12221, 2019. Available: <https://arxiv.org/abs/1903.12221> Accessed: Jan. 1, 2025.
- [19] G. McGrath and P. R. Brenner, “Serverless Computing: Design, Implementation, and Performance,” in *2017 IEEE 37th International Conference on Distributed Computing Systems Workshops (ICDCSW)*, pp. 405–410, 2017, <https://doi.org/10.1109/ICDCSW.2017.36>
- [20] A. Mohan et al., “Agile Cold Starts for Scalable Serverless,” in *11th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 19)*, 2019. Available: <https://www.usenix.org/conference/hotcloud19/presentation/mohan> Accessed: Jan. 1, 2025.
- [21] E. Oakes et al., “Rapid Task Provisioning with Serverless-Optimized Containers,” in *USENIX Annual Technical Conference (ATC 18)*, pp. 57–70, 2018. Available: <https://www.usenix.org/conference/atc18/presentation/oakes> Accessed: Jan. 1, 2025.
- [22] G. Persson and W. Branth Sjöberg, “Mitigating Serverless Cold Starts Through Predicting Computational Resource Demand: Predicting Function Invocations Based on Real-Time User Navigation,” bachelor’s thesis, Jönköping University, 2023.
- [23] S. Rajbhandari et al., “ZeRO-Infinity: Breaking the GPU Memory Wall for Extreme Scale Deep Learning,” in *SC21: International Conference for High Performance Computing, Networking, Storage and Analysis*, 2021, <https://doi.org/10.1145/3458817.3476205>
- [24] Sedai, “AWS Autonomous Concurrency,” 2025. Available: <https://www.sedai.io/blog/the-end-of-cold-starts-autonomous-concurrency-for-aws-lambda> Accessed: Jan. 1, 2025.
- [25] Y. Sheng et al., “FlexGen: High-Throughput Generative Inference of Large Language Models with a Single GPU,” arXiv:2303.06865, 2023. Available: <https://arxiv.org/abs/2303.06865> Accessed: Jan. 1, 2025.
- [26] S. Shillaker and P. R. Pietzuch, “A Provider-Friendly Serverless Framework for Latency-Critical Applications,” 2018. Available: <https://api.semanticscholar.org/CorpusID:51997872>
- [27] SlideShare, “Become a Serverless Black Belt: Optimizing Your Serverless Applications,” 2025. Available: <https://www.slideshare.net/slideshow/become-a-serverless-black-belt-optimizing-your-serverless-applications-srv401-reinvent-2017/840310167> Accessed: Jan. 1, 2025.
- [28] Snowflake, “LLM Inference Guide,” 2025. Available: <https://www.snowflake.com/guides/llm-inference/> Accessed: Jan. 1, 2025.
- [29] P. Vahidinia et al., “Cold Start in Serverless Computing: Current Trends and Mitigation Strategies,” in *2020 International Conference on Omni-layer Intelligent Systems (COINS)*, pp. 1–7, 2020, <https://doi.org/10.1109/COINS49042.2020.9191377>
- [30] Y. Wang et al., “Tabi: An Efficient Multi-Level Inference System for Large Language Models,” in *Proceedings of the 18th European Conference on Computer Systems (EuroSys ’23)*, pp. 233–248, 2023, <https://doi.org/10.1145/3552326.3587438>
- [31] J. Wen, Z. Chen, X. Jin, and X. Liu, “Rise of the Planet of Serverless Computing: A Systematic Review,” *ACM Transactions on Software Engineering and Methodology*, vol. 32, no. 5, pp. 1–61, 2023, <https://doi.org/10.1145/3579643>
- [32] Wikipedia, “Serverless Computing,” 2025. Available: https://en.wikipedia.org/wiki/Serverless_computing Accessed: Jan. 1, 2025.
- [33] B. Wu et al., “Fast Distributed Inference Serving for Large Language Models,” arXiv:2305.05920, 2023. Available: <https://arxiv.org/abs/2305.05920> Accessed: Jan. 1, 2025.
- [34] H. Yu, R. Basu Roy, et al., “RainbowCake: Mitigating Cold-starts in Serverless with Layer-wise Container Caching and Sharing,” in *ASPLOS ’24*, 2024, <https://doi.org/10.1145/3617232.3624871>
- [35] K. Zhao, Z. Zhou, L. Jiao, S. Cai, F. Xu, and X. Chen, “Taming Serverless Cold Start of Cloud Model Inference With Edge Computing,” *IEEE Transactions on Mobile Computing*, vol. 23, no. 8, pp. 8111–8128, 2024, <https://doi.org/10.1109/TMC.2023.3348165>