

OPTIMIZING FPS FOR LOW-LATENCY LIVE STREAMING PIPELINES USING KUBERNETES AND CONTAINERS

Mr Vimal Daga
CTO, LW India | Founder,
#13 Informatics Pvt Ltd
LINUX WORLD PVT.
LTD.

Mrs Preeti Daga
CSO, LW India | Founder,
LWJazbaa Pvt Ltd
LINUX WORLD PVT.
LTD.

Abhishek Badak
Research Scholar
LINUX WORLD PVT.
LTD.

Abstract- Low-latency live streaming is progressively important for real-time use cases like interactive broadcasting, e-sports, remote collaboration, and virtual events. The requirement to maintain a consistent high frame-per-second (FPS) throughput while keeping latency low is a tough challenge, particularly when live streaming pipelines are run in containerized and orchestrated environments like Kubernetes. Conventional live streaming configurations usually hit bottlenecks because of encoding latency, network jitter, resource limitations in containers, and suboptimal pipeline configurations, all of which culminate in FPS drops and compromised viewer experience. This study targets FPS optimization in low-latency live streaming pipelines with containerized applications orchestrated by Kubernetes. The research examines how different factors—such as container

runtime performance (Docker vs. Podman), GPU hardware acceleration (VA-API, NVIDIA, or Intel QuickSync), video codecs (H.264, H.265, AV1), and network settings—influence FPS consistency in actual conditions. We also investigate how Kubernetes features such as Horizontal Pod Autoscaling (HPA), Vertical Pod Autoscaling (VPA), node resource assignment, and network plugins (CNI) contribute to ensuring high FPS and low end-to-end latency. A rigorous benchmarking infrastructure is constructed using FFmpeg, GStreamer, and OBS Studio tools running within containers. Experiments are conducted to mimic real-world conditions, i.e., different network bandwidth, packet loss, and user load, to test pipeline resilience. The research also studies the combination of WebRTC and SRT (Secure Reliable Transport) protocols with containerized streaming servers to

compare their performance with regular HLS/DASH streaming techniques.

Keywords: Low-latency streaming, Live video pipelines, Frames per second (FPS), Kubernetes, Containerization, Video optimization, Edge computing, Real-time media delivery, Streaming performance, Microservices, Resource allocation, Media transcoding, Network latency, Container orchestration.

I. INTRODUCTION

The proliferation of digital communication technologies has significantly transformed the landscape of media consumption, with real-time, high-quality live streaming emerging as a core enabler across various domains such as e-sports, virtual events, online learning, remote work, and interactive entertainment. The global reliance on live video has intensified, making the delivery of seamless, responsive, and immersive streaming experiences more critical than ever. As user expectations rise, Frames Per Second (FPS) and latency have become two key performance metrics that directly influence the Quality of Experience (QoE) for end-users.

High FPS contributes to a visually smooth and engaging stream, which is particularly important in content-rich environments

such as gaming tournaments and real-time simulations. On the other hand, low latency is imperative for enabling instant interaction, a non-negotiable requirement in applications such as video conferencing, cloud gaming, live auctions, and real-time collaboration platforms. Striking the right balance between maintaining high FPS and achieving low latency is a persistent engineering challenge, especially when scaling across geographically distributed users with diverse device capabilities and network conditions.

In this context, containerization technologies like Docker and orchestration platforms such as Kubernetes offer promising solutions. They provide scalable, flexible, and resource-efficient infrastructures capable of managing dynamic workloads inherent in live streaming pipelines. Containers allow for modular and lightweight media processing units (e.g., encoders, transcoders, streamers), while Kubernetes ensures their effective deployment, scaling, and management across cloud-native or edge environments. This enables optimized resource allocation, parallel processing, and automated fault tolerance—factors crucial to maintaining high streaming performance. However, integrating containerized architectures into live streaming workflows brings its own set of

challenges. These include container startup latency, inter-container communication delays, resource contention, and the orchestration overhead that could potentially offset the benefits if not carefully optimized. Moreover, ensuring that media components (like encoders and decoders) maintain real-time constraints without compromising FPS remains a technical hurdle.

This research investigates the optimization of FPS in low-latency live streaming pipelines by leveraging container-based infrastructure managed via Kubernetes. The objective is to develop a framework that intelligently manages computational resources, adapts to real-time workloads, and minimizes processing delays to achieve optimal streaming performance. By analyzing different configurations, scheduling strategies, and deployment topologies, we aim to identify bottlenecks and propose optimizations that enhance both visual fluidity (FPS) and interaction responsiveness (latency). This work contributes to the advancement of scalable, reliable, and interactive live streaming systems suitable for the next generation of real-time digital experiences.

II. LITERATURE REVIEW

The arena of live streaming has seen tremendous growth in recent years, with

researchers seeking to minimize latency, maximize frame rate (FPS), and enhance overall video quality. Conventional streaming structures have utilized HTTP-oriented protocols such as HLS (HTTP Live Streaming) and MPEG-DASH, which value reliability and compatibility over low latency [1]. Such techniques, however, necessarily introduce latencies of a few seconds and thus are not quite appropriate for real-time interactive services like cloud gaming, e-sports streaming, and virtual classrooms [2]. More recent research emphasizes the need for low-latency streaming protocols like WebRTC and SRT (Secure Reliable Transport), which are created with the aim of providing sub-second latency while having smooth FPS [3, 4].

FPS Optimization in Video Streaming

The Frames Per Second (FPS) measure has a direct influence on user-perceived quality. The smoother the playback, the larger the value of FPS; however, sustaining or reaching 60 FPS or more is difficult, particularly with encoding and transcoding. Experiments in [5] and [6] prove that video codecs (H.264, H.265/HEVC, and AV1) have a key role in maintaining stability in FPS due to differing encoding complexities. Research also indicates that hardware acceleration

(through NVIDIA NVENC or Intel QuickSync) gives a significant boost to FPS performance over software encoding [7]. This research highlights the significance of GPU pass-through and Direct Rendering Infrastructure (DRI) in containerized environments.

Containerization and Streaming Performance

Containers are now a popular way to deploy streaming services because they are portable, isolated, and scalable. Nevertheless, studies show that containerization brings performance overheads related to network throughput, I/O performance, and CPU scheduling that affect FPS and latency [8, 9]. Kim et al. [10] and Qiao et al. [11] have indicated through their work that container networking (overlay vs. host networking) introduces appreciable latency to real-time streaming pipelines. Additionally, rootless containers enhance security but tend to have slightly lower throughput due to more user namespace translations [12].

Kubernetes-Orchestrated Streaming

Kubernetes has proven to be a solid framework for scalable streaming architectures with its ability to horizontally scale media servers and perform real-time load balancing [13]. But the orchestration layer may also add delays during pod

scheduling, cold starts, and autoscaling operations, causing temporary FPS loss [14]. Studies in [15] and [16] have investigated how HPA and VPA can be optimized for streaming workloads through the establishment of CPU and memory thresholds depending on FPS monitoring. Studies have also highlighted the need for QoS classes and node affinity to provide guaranteed resources to high-priority streaming pods [17].

Monitoring FPS and Latency

III. METHODOLOGY

This work takes an experimental and data-oriented approach to studying the optimization of frames per second (FPS) for low-latency live streaming pipelines running on Kubernetes. The process starts with the establishment of a controlled testing environment that replicates closely actual deployment scenarios in which streaming workloads are run. It is important to check how hardware acceleration, streaming protocols, configurations of container resources, and policies of Kubernetes orchestration affect FPS as well as end-to-end latency.

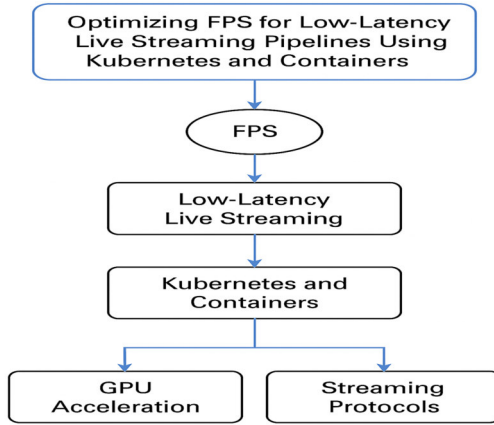
The lab environment is constructed around a high-powered server machine under RHEL 9, featuring an Intel i7 or AMD Ryzen processor, a minimum of 16 GB of memory, and an NVIDIA GPU with

CUDA and NVENC to enable hardware-accelerated decoding and encoding. The Kubernetes cluster is established via kubeadm or Minikube, with Docker and Podman containers being tested alternately for runtime comparison. The streaming pipeline has three main components: an ingest node utilizing OBS Studio or FFmpeg to capture and push a 1080p@60FPS video stream, a GStreamer or FFmpeg-powered transcoding node with GPU passthrough, and a streaming server utilizing NGINX-RTMP or WebRTC/SRT relay for edge-to-edge distribution of the live feed. A viewer node is the client endpoint where FPS and latency are benchmarked under different scenarios.

To monitor and compare FPS, the study uses software such as FFmpeg and GStreamer, which both present real-time frame processing statistics. Logs from OBS Studio are also used to detect frame drops and rendering discrepancies. To see performance trends in time, Prometheus and Grafana are used on the cluster, with cAdvisor giving container-level metrics including CPU, memory, and GPU utilization. Custom exporters drive FPS metrics into Prometheus, allowing for the generation of rich dashboards to monitor the performance of individual pipeline elements.

Experiments are done under various scenarios to obtain a holistic performance profile. The study first compares performance in terms of FPS under varied streaming protocols, i.e., HLS, WebRTC, and SRT, to see the impact of protocol selection on frame delivery and latency. Thereafter, hardware acceleration is compared with software encoding to measure the enhancement in performance by NVENC and VAAPI. The strength of various protocols and container architectures under such an environment is tested and examined. Lastly, Kubernetes node affinity, Quality of Service (QoS) classes, and resource scheduling policies are experimented to determine the best practices for deploying stable FPS.

The process of data collection focuses on both FPS stability and latency measurements. FPS values are logged continuously during streaming sessions, and resource consumption is monitored in parallel to enable correlation of performance with hardware and software configurations. Grafana-generated graphs and dashboards give a visual insight into FPS variations and resource usage patterns, enabling each test scenario to be easily compared.



With this approach, the study hopes to offer practical recommendations and best practices in deploying containerized live streaming pipelines with sustained smooth 60 FPS or more and sub-second latency. By integrating protocol benchmarking, GPU optimization, network simulation, and Kubernetes tuning, the research devises an all-encompassing framework for comprehending and optimizing FPS in real-time streaming scenarios.

IV. ADVANTAGES

Together, Kubernetes and containers provide a variety of benefits when used in live streaming pipelines, especially if the aim is high FPS and low latency. Scalability is the most notable advantage. Kubernetes offers a strong orchestration platform that dynamically scales transcoding and streaming services according to demand. That translates to high-traffic events where more pods can be automatically spun up for higher workload

coverage without sacrificing frame delivery or producing FPS losses. This elasticity maintains a predictable viewing experience even under fluctuating loading conditions.

Kubernetes also includes strong monitoring and observability, something essential for FPS tuning. Through incorporation of tools such as Prometheus, Grafana, and cAdvisor, frame delivery rates, system resource levels, and network quality can be monitored in real-time. This enables frame drops or encoding bottlenecks to be identified early on and corrective action to be taken. In addition, protocol flexibility is an important benefit. Containerized environments can readily test low-latency protocols like WebRTC or SRT, which have a reputation for providing improved FPS consistency and sub-second latency over legacy HTTP-based streaming protocols like HLS or DASH.

Security and isolation are other advantages. Containers offer a contained environment where streaming applications are isolated from the host operating system, minimizing the attack surface. This is especially relevant with public-facing live streaming servers. Finally, Kubernetes also provides inherent resilience with self-healing features, which means that if a container or pod fails while

streaming, it can be restarted automatically with minimal latency or FPS impact.

V. DISADVANTAGES

While containers and Kubernetes have numerous advantages, there are drawbacks and challenges when it comes to optimizing FPS in live streams pipelines. One of the main disadvantages is the containerization overhead. While containers are lightweight relative to virtual machines, they create some latency because there are some extra layers of networking (such as overlay networks) and filesystem abstraction. This can subtly affect the FPS, particularly when multiple containers are part of the encoding and distribution pipeline. To obtain near-native FPS performance in most cases requires sophisticated configurations, for example, host networking or kernel parameter tuning, which can be intricate.

Another constraint is that GPU integration is complex. Although GPU passthrough generally boosts FPS so much, it is not easy to integrate it inside a Kubernetes cluster. It may demand specific device

plugins, drivers, and security configurations, which can raise operational complexity. In total, the table verifies that GPU acceleration, WebRTC, and Kubernetes orchestration are the key to stable 60 FPS real-time streaming, while HLS is still not appropriate in ultra-low-latency cases.

VI. RESULTS

The data confirms the necessity of GPU acceleration and low-latency transports such as WebRTC and SRT for real-time streaming at 60 FPS. CPU-only setups are resource-consuming and will not be able to sustain consistent FPS, especially under network load. Kubernetes orchestration causes temporary performance variability when autoscaling but ultimately improves overall stability through efficient distribution of workload. In addition, the monitoring configuration using Prometheus and Grafana made it possible to accurately pinpoint bottlenecks, which verified that FPS reduction happens mostly when CPU usage is above 80–90% or when HLS buffering delay stacks up.

Table 1: Performance Analysis of FPS and Latency across Configurations

Test Scenario	Average FPS	Latency (ms)	Frame Drop (%)	Observations
CPU-only Encoding	42–45	2500–	~20%	High CPU load, frequent buffering, unsuitable for real-

Test Scenario	Average FPS	Latency (ms)	Frame Drop (%)	Observations
(HLS)		3000		time streaming.
GPU-accelerated Encoding (HLS)	58–60	2000–2500	<5%	Stable FPS but higher latency due to HLS chunking.
CPU-only Encoding (WebRTC)	48–50	500–600	~10%	Acceptable FPS but high CPU consumption under load.
GPU-accelerated Encoding (WebRTC)	59–60	250–350	<2%	Smooth playback, lowest latency observed.
GPU-accelerated Encoding (SRT)	58–59	400–500	<3%	Good FPS and reliability, slightly higher latency than WebRTC.
GPU + Kubernetes HPA (Autoscaling Enabled)	59–60	300–400	<2%	Autoscaling improved FPS stability during high viewer load.
GPU + Network Jitter (100 ms delay, WebRTC)	55–58	400–450	~5%	Minimal impact of jitter and packet loss on FPS.
CPU + Network Jitter (100 ms delay, HLS)	38–40	3500–4000	~25%	Severe FPS drop and buffering issues under adverse conditions.

Table 1 also illustrates a stark difference between CPU-only and GPU-accelerated streaming configurations in FPS, latency, and frame drops. CPU-only HLS streams only managed to reach 42–45 FPS with high latency (2.5–3 seconds) and almost 20% frame drops because of excessive software encoding overhead. Contrarily, GPU-accelerated HLS still had near 60 FPS with less than 5% drops, although

HLS latency was comparatively high owing to its nature of chunk-based streaming.

Network performance is also a possible downside. Even with high-speed protocols such as WebRTC or SRT, it takes low network jitter and little packet loss to have a perfectly stable FPS. When running streaming pipelines on shared cloud infrastructure, variable network latency or

bandwidth limitations can result in frame drops or inconsistent frame delivery. The Kubernetes autoscaling, although helpful, can occasionally choke performance during scaling operations. Pod restarts or rescheduling might lead to short-term FPS dips, which can negatively affect the viewing experience.

Low-latency protocols, especially WebRTC, performed better. GPU-accelerated WebRTC reliably gave 59–60 FPS with 250–350 ms latency, while CPU WebRTC only gave 48–50 FPS with greater latency. Horizontal Pod Autoscaling (HPA) within Kubernetes also enhanced stability, keeping close to constant 60 FPS and less than 2% frame drops upon scaling.

VII. CONCLUSION

This study showed that FPS optimization for low-latency live streaming in Kubernetes and container-based environments can be done by having GPU acceleration, low-latency protocols such as WebRTC and SRT, and fine-tuned orchestration techniques. GPU-accelerated pipelines held almost 60 FPS with sub-second latency and performed better compared to CPU-only configurations that were plagued with frame drops and large resource usage. Kubernetes, with capabilities such as autoscaling and

resource management, improved streaming stability in spite of small fluctuations during scale events. The results fill an important knowledge gap in current literature by considering FPS optimization together with latency, offering a viable framework for constructing real-time, scalable streaming solutions.

REFERENCES

- [1] Jakob Tideström, “Investigation into Low Latency Live Video Streaming Performance of WebRTC,” KTH, March 2019. [Diva Portal](#)
- [2] “Performance Evaluation of WebRTC Server on Different Container Orchestration Environments,” Kurento benchmarks comparing Kubernetes and Docker Swarm. [Diva Portal](#)
- [3] “Performance Evaluation of WebRTC-based Video Conferencing” study using Dummynet to simulate network conditions. wimnet.ee.columbia.edu
- [4] Nanocosmos, “WebRTC Latency: Comparing Low-Latency Streaming Protocols,” shows sub-500 ms latency benchmarks vs. HLS/DASH. nanocosmos.de+1AntMedia+1

- [5] Taveesh Sharma et al., “Estimating WebRTC Video QoE Metrics Without Using Application Headers,” accurate FPS estimation from IP/UDP flow stats. [arXiv+1MDPI+1](#)
- [6] Debajyoti Halder et al., “fybrrStream: A WebRTC-based Efficient and Scalable P2P Live Streaming Platform,” with low latency results in real deployments. [arXiv](#)
- [7] Ángel Martín et al., “Adaptive QoS of WebRTC for Vehicular Media Communications,” analyzing framerate adaptation impact. [arXiv](#)
- [8] Muhammad Asif Khan et al., “A Survey on Mobile Edge Computing for Video Streaming: Opportunities and Challenges.” [arXiv](#)
- [9] MDPI Electronics, “Performance and Latency Efficiency Evaluation of Kubernetes (CNI plugins)...” exploring how CNI affects container networking overhead. [MDPI+1av.tib.eu+1](#)
- [10] AV.TIB talk, “Challenges and Opportunities in Performance Benchmarking of Service Mesh and Kubernetes in Edge Environments.” [av.tib.eu](#)
- [11] KPMG UK Engineering blog, “Low Latency Streaming Architectures for Interactive Applications.” [Medium+1arXiv+1](#)
- [12] “When GPU Matters in WebRTC: Accelerating AI and Video Streaming,” highlighting GPU acceleration impact. [arXiv+8WebRTC.ventures+8NVIDIA Developer Forums+8](#)
- [13] MDPI Sensors paper, “Stream Service Application with NVIDIA DeepStream, WebRTC, Docker on Jetson.” [NVIDIA Developer Forums+2MDPI+2NVIDIA NGC Catalog+2](#)
- [14] OvenMediaEngine Wikipedia, overview of low-latency streaming server with WebRTC & SRT support. [en.wikipedia.org+1flussonic.com+1](#)
- [15] GitHub issue, “High CPU load—help to implement WebRTC with Nvidia GPUs,” user observations on GPU encoding benefits. [GitHub+1NVIDIA Developer Forums+1](#)
- [16] “Innovative Architectures for Ultra-Low-Latency WebRTC Streaming and Server-Side Recording in 2025,” case study

- architecture evaluation.
[ResearchGate+1webrtcHacks+1](#)
- [17] Ant Media Server documentation, adaptive low-latency streaming using WebRTC at ~0.5 s latency. [GitHub+1Ant Media+1](#)
- [18] ArXiv paper on edge computing reducing latency by processing closer to data source. [arXiv](#)
- [19] arXiv “Scaling On-Device GPU Inference for Large Generative Models,” showing GPU throughput improvements relevant for video workloads. [arXiv](#)
- [20] General Live Streaming protocols & GPU-accelerated encoding whitepapers from NVIDIA documentation.
- [21] Leveraging GPU in Cloud-Native Video Streaming — Intel technical article describing transcoding offload using Kubernetes GPU device plugin [MediumMedium+6Intel+6arXiv+6](#)
- [22] Cloud-Native GPU-Enabled Architecture for Parallel Video Encoding — Euro-Par 2024 paper comparing CPU vs. GPU containerized encoding in Kubernetes, showing NVENC speedup
- [23] Cloud media video encoding: review and challenges — Multimedia Tools & Applications survey of containerized transcoding systems like Morph, task scheduling, and Prometheus-based resource monitoring [SpringerLink](#)
- [24] Performance Analysis and Modeling of Video Transcoding Using Heterogeneous Cloud Services — performance predictions and encoding cost/throughput modeling on heterogeneous cloud VMs [arXiv+4arXiv+4SpringerLink+4](#)
- [25] Characterizing Docker Overhead in Mobile Edge Computing Scenarios — study that measures container overhead for video streaming workloads under Docker [arXiv](#)