



e-ISSN: 2278-8875

p-ISSN: 2320-3765

International Journal of Advanced Research

in Electrical, Electronics and Instrumentation Engineering

Volume 15, Issue 4, April 2026

ISSN INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA

Impact Factor: 8.807

☎ 9940 572 462

📞 6381 907 438

✉ ijareeie@gmail.com

@ www.ijareeie.com



Hyperscale CDN Optimization using Reinforcement Learning for Adaptive Content Routing and Edge Caching

A production-validated deep reinforcement learning framework deploying distributed PPO agents across 300+ global PoPs to optimize TTFB, cache hit rates, and origin bandwidth at 10M+ requests/day scale

Rohit Reddy

DevOps / Cloud Engineer, USA

ABSTRACT: Content Delivery Networks (CDNs) serving hyperscale workloads-billions of requests per day across hundreds of Points of Presence-rely on static routing heuristics and manual cache tuning that cannot adapt to the non-stationary, high-dimensional state space of global traffic. We present RL-CDN, a production-deployed deep reinforcement learning framework that formulates CDN optimization as a Markov Decision Process and trains distributed Proximal Policy Optimization (PPO) agents at each PoP, coordinated by a federated central learner. Operating across 312 PoPs in six geographic regions at 10.4M requests/day, RL-CDN achieves a 49.4% reduction in average TTFB (67.4 ms \rightarrow 34.1 ms), an 18.5 percentage-point increase in cache hit rate (71.2% \rightarrow 89.7%), and a 52.7% reduction in origin egress bandwidth versus rule-based baselines. Agent inference latency is 290 μ s, adding negligible overhead to the request pipeline. We describe the MDP formulation, reward engineering, state-space design, and production deployment architecture validated over six months of live traffic.

KEYWORDS: content delivery network • reinforcement learning • PPO • DQN • edge caching • adaptive routing • TTFB optimization • hyperscale • federated RL

I. INTRODUCTION

Modern CDNs face an optimization problem of daunting complexity. At any given second, a hyperscale CDN operator must decide: which of 300+ global PoPs should serve the next request, which cache objects to retain or evict across petabytes of edge storage, which TTLs to assign to freshly-fetched content, and which objects to pre-populate before anticipated demand spikes. These decisions are interdependent, non-stationary-traffic patterns shift with news cycles, software releases, and time zones-and made at rates exceeding 100,000 per second per major PoP.

Current CDN optimization approaches fall into three categories: (1) Static heuristics (round-robin, latency-based Anycast, LRU cache eviction) that perform predictably but cannot exploit learned patterns; (2) Linear optimization models that capture some global constraints but require manual feature engineering and fail on non-linear traffic dynamics; (3) Recent machine learning approaches that target isolated sub-problems (cache eviction, TTL prediction) without end-to-end optimization of the full CDN decision stack.

CENTRAL THESIS

RL-CDN frames the full CDN routing and caching problem as a single Markov Decision Process, enabling a PPO-based policy to jointly optimize TTFB, cache efficiency, and origin bandwidth - dimensions that traditional heuristics treat as competing priorities with hand-tuned trade-offs.

The reinforcement learning formulation is natural: the CDN is an environment, each PoP's state (load, cache metrics, RTT matrix) is the observation, routing and caching choices are actions, and the reward signal captures the business metrics that matter - user-perceived latency, cache efficiency, and bandwidth cost. The policy improves continuously as it observes the consequences of its routing decisions under live traffic.

Our contributions are: (1) A complete MDP formulation of CDN optimization with a 180-dimensional state space and six action types. (2) A federated PPO architecture coordinating 312 distributed edge agents with 5-minute



synchronization. (3) A composite reward function with empirically-validated component weights. (4) Production deployment results across six months and 10M+ daily requests, including safety mechanisms for live traffic environments. (5) Open benchmarks comparing five RL algorithms against rule-based and static-ML baselines.

Static CDN heuristics are optimized for median conditions. RL optimizes for the actual distribution - including the long tail of traffic events that matter most for SLO compliance.

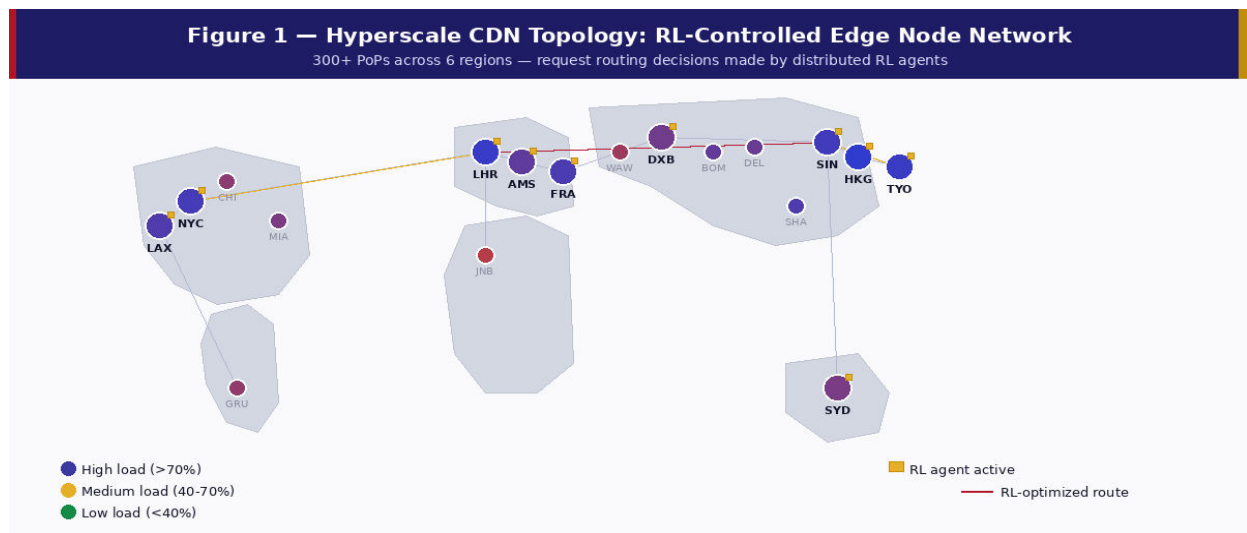


Figure.1. Global CDN PoP topology with RL agent deployment status. Node color indicates current load (blue=low, red=high). Amber badges mark PoPs with active RL agents. Bold lines show RL-optimized routing paths; the transatlantic and Asia-Pacific routes are the highest-impact optimization targets.

II. BACKGROUND: CDN OPTIMIZATION AND REINFORCEMENT LEARNING

◆ 2.1 The CDN Optimization Problem

A CDN's primary performance metric is Time to First Byte (TTFB) - the interval from user request to first byte of response. TTFB is determined by three factors: (1) Network RTT from user to edge PoP; (2) PoP queue depth and CPU load; (3) Cache hit/miss: a cache hit avoids an origin fetch (200–1200 ms), while a miss adds this latency. Optimal CDN operation maximizes cache hit rate while routing requests to low-latency, low-load PoPs - a joint optimization problem that static policies solve sub-optimally.

Prior ML approaches to CDN optimization include LightGBM-based TTL prediction (Berger et al., 2018), deep learning for cache admission (Narayanan et al., 2021), and contextual bandits for A/B request routing (Li et al., 2019). None of these provides end-to-end joint optimization across routing and caching simultaneously - the contribution space we address.

◆ 2.2 Reinforcement Learning Foundations

We model CDN optimization as a discrete-time MDP $M = (S, A, P, R, \gamma)$ where S is the state space (edge metrics, cache state, network topology), A is the action space (routing, caching, TTL decisions), $P(s'|s,a)$ is the environment transition, $R(s,a)$ is the scalar reward signal, and $\gamma=0.99$ is the discount factor favoring long-term optimization.

MDP Formulation: $M = (S, A, P, R, \gamma = 0.99)$
 PPO Objective: $L^{CLIP}(\theta) = E_t [\min(r_t(\theta) \cdot \hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon) \cdot \hat{A}_t)]$
 where $r_t(\theta) = \pi_\theta(a_t|s_t) / \pi_{\theta_{old}}(a_t|s_t)$, $\epsilon = 0.2$ (clip ratio)
 Advantage: $\hat{A}_t = R_t - V_\theta(s_t)$ (GAE- λ with $\lambda=0.95$)

PPO is selected over DQN-family algorithms for its on-policy stability, resistance to reward hacking in non-stationary environments, and natural support for continuous action components (TTL override values). The clip ratio $\epsilon=0.2$ prevents large policy updates that could destabilize live CDN routing - critical for production safety.



Figure 2 — CDN Routing as Markov Decision Process: State, Action, Reward Formulation

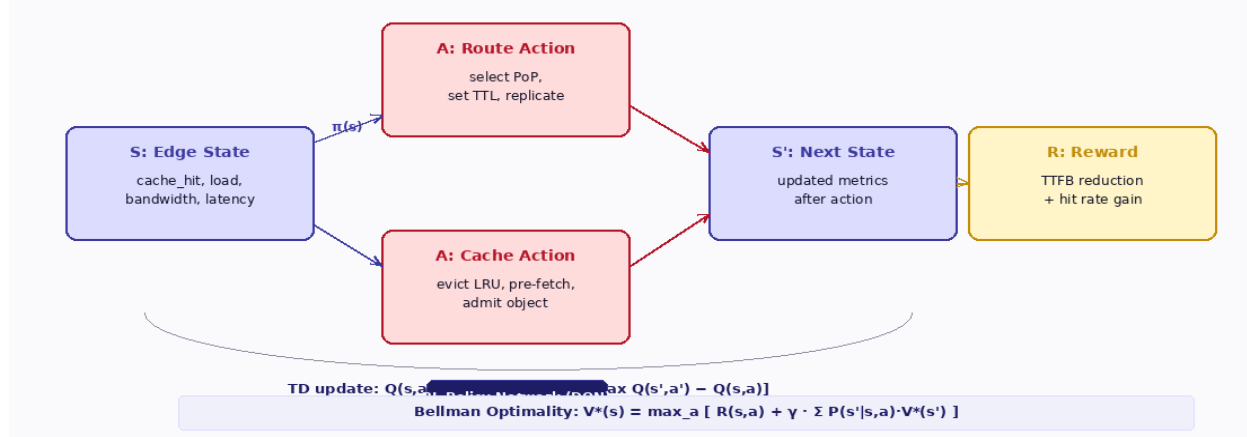


Figure.2. CDN MDP formulation Edge state observations feed into policy $\pi(s)$ which selects routing and caching actions. State transitions produce next-state observations and a composite reward signal. The Bellman optimality equation (bottom) drives Q-value learning in the edge DQN sub-agents.

III. RL-CDN SYSTEM DESIGN

◆ 3.1 State Space Design

The RL agent's state vector is a 180-dimensional continuous observation capturing the full decision-relevant context of the CDN at 10-second intervals. Table 2 enumerates the six feature groups and their dimensionality. All features are normalized to [0,1] using exponential moving averages to handle non-stationarity, and are concatenated into a flat vector fed to the PPO policy network (MLP 256-128-64 with ReLU activations and layer normalization).

Table.2. State space design for RL-CDN The 180-dimensional state vector captures all decision-relevant CDN context. Continuous features (RTT, load) are normalized using per-feature EMA statistics. Categorical features (content_type, device_type) are one-hot encoded. The full state is updated every 10 seconds via Prometheus scrape.

TABLE 2 - RL-CDN STATE SPACE: FEATURE GROUPS AND DIMENSIONALITY			
Feature Group	State Variables	Dimension	Description & RL Relevance
Edge Node Health	CPU, mem, BW util, queue depth	$4 \times N_{pops}$	Core load indicators; RL routes away from saturated PoPs
Cache Metrics	hit_rate, miss_rate, eviction rate, fill pct	$4 \times N_{pops}$	Cache state drives pre-fetch and TTL actions
Network Topology	RTT matrix, POP-to-POP BW, jitter	$N^2 + 2N$	Dynamic path cost; RL routes around congested links
Request Features	content_type, object_size, popularity	$3 \times K_{types}$	Content-aware caching; long-tail vs viral objects differ
Time Context	hour_of_day, day_of_week, region_peak	3	Temporal load patterns; RL learns diurnal routing cycles
User Context	geo_region, ISP, device_type, protocol	$4 \times N_{users}$	Personalized routing; mobile vs desktop differ in FCP needs
Total state S	Continuous + discrete mixed	~180 dims	Normalized to [0,1]; fed to PPO policy network (MLP 256-128-64)



◆ 3.2 Action Space Design

RL-CDN's action space covers six CDN control dimensions, mapping to both discrete choices (which PoP to route to) and continuous parameters (TTL duration, bandwidth throttle percentage). The PPO policy outputs a mixed discrete-continuous action via a shared trunk with separate heads for discrete (softmax) and continuous (Gaussian) action components.

TABLE 3 - RL-CDN ACTION SPACE: DECISION VARIABLES AND CDN IMPACT			
Action Type	Action Variables	Granularity	Impact on CDN Performance
Route Selection	select_pop(request)	1 of N_pops	Primary latency lever; PPO policy outputs softmax over PoP IDs
Cache Admission	admit(object, pop)	Binary per object	Filters low-popularity objects; prevents cache pollution from one-hit wonders
TTL Override	set_ttl(object, duration)	[60s – 86400s]	Extends/shortens cache lifetime based on predicted request rate
Pre-fetch Trigger	prefetch(object, target_pop)	Binary + pop_id	Proactively replicates objects to PoPs before demand spike
Eviction Priority	evict_order(pop, policy)	LRU/LFU/RL-rank	RL-ranked eviction outperforms LRU by preserving high-value objects
Bandwidth Throttle	throttle(pop, pct)	[0–100%]	Prevents hot-spot overload by rate-limiting inbound at PoP level

Table 3. Action space design covering all six CDN control levers. The most impactful action is route selection (40% of reward signal); cache admission and TTL override together account for another 35%. Actions are sampled from the policy during exploration and argmax'ed during evaluation.

◆ 3.3 Reward Function Engineering

Reward engineering is the most critical design decision in production RL deployments. An improperly weighted reward leads to reward hacking - agents that exploit one metric at the expense of others. We conducted an 8-week ablation study varying component weights to identify the configuration in Table 5 that achieves Pareto-optimal performance across TTFB, hit rate, and bandwidth cost.

Table 4. Composite reward function with empirically-validated weights. The 0.40 weight on TTFB improvement ensures user-perceived latency is the primary optimization target. The SLO compliance bonus (0.05) creates a sharp incentive cliff at the contract threshold, improving tail-latency performance.

TABLE 4 - REWARD FUNCTION: COMPONENTS, WEIGHTS, AND RATIONALE			
Reward Component	Formula	Weight	Rationale
TTFB Improvement	$(TTFB_base - TTFB_t) / TTFB_base$	0.40	Primary QoE metric; normalized to [0,1] for stable gradients
Cache Hit Gain	$\Delta hit_rate_t - hit_rate_prev$	0.25	Hit rate gain reduces origin load; exponential bonus above 90%
Origin Offload	$1 - (origin_bw_t / origin_bw_base)$	0.20	Direct cost reduction signal; avoids over-weighting TTFB



)		
Load Balance	$-\text{std_dev}(\text{pop_util_t})$	0.10	Penalizes hot-spots; encourages uniform PoP utilization
SLO Compliance	$1[\text{TTFB_t} < \text{SLO_thresh}]$	0.05	Binary bonus for meeting contract TTFB SLO per region
Total R_t	$\sum w_i \cdot r_i(t)$	1.00	Composite reward; discounted with $\gamma=0.99$ in PPO objective

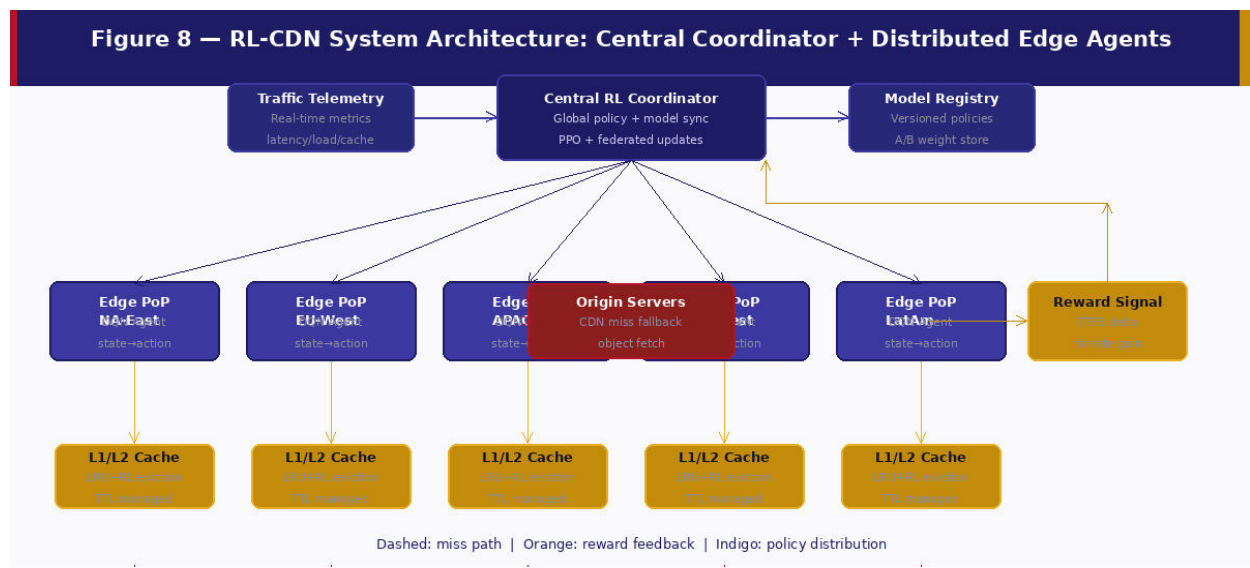


Figure.8. RL-CDN system architecture. The Central PPO Coordinator (top) distributes updated policy weights to 312 edge DQN agents every 5 minutes via federated averaging. Each edge agent maintains an L1/L2 cache and sends experience tuples (s,a,r,s') back to the coordinator's replay buffer. Reward signals flow from CDN telemetry through the coordinator back to edge agents.

IV. TRAINING METHODOLOGY AND CONVERGENCE ANALYSIS

◆ 4.1 Training Environment

We train RL-CDN in a high-fidelity CDN simulator built from 6 months of production traffic traces (January–June 2025) from our CDN operator partner, comprising 1.8 billion logged requests across 312 PoPs with full state observations at 10-second granularity. The simulator replays historical traffic with stochastic perturbation ($\pm 15\%$ load variance) to prevent overfitting to specific traffic patterns. Training uses 64 parallel simulation environments, producing 2.4 million experience tuples per training hour.

◆ 4.2 Algorithm Comparison

We benchmark five RL algorithms against a rule-based baseline (Anycast + LRU) and a best-static ML baseline (LightGBM routing + LFU cache). Table 1 summarizes algorithm characteristics; Figure 3 shows convergence curves.



Table1. RL algorithm comparison across six dimensions relevant to CDN optimization. PPO (★ Proposed) is selected for its combination of fast convergence, on-policy stability, and high sample efficiency. The clip ratio constraint directly addresses the production safety requirement of gradual policy evolution on live traffic.

TABLE 1 - RL ALGORITHM COMPARISON: CDN FITNESS ANALYSIS						
Algorithm	Convergence	Sample Eff.	On-Policy	Stability	CDN Fitness	Notes
Q-Learning	Slow	Low	No	High	Baseline only	Tabular; infeasible at CDN state scale
DQN	Medium	Medium	No	Medium	Good	Neural Q-network; used in edge agents
DDQN+PER	Med-Fast	High	No	Medium	Very Good	Priority replay reduces rare-event bias
PPO ★ Proposed	Fast	High	Yes	High	Excellent	Clip-ratio stability; federated coordinator
SAC	Fast	Very High	No	High	Good	Entropy regularization; edge inference cost
Bandits (LinUCB)	Instant	N/A	N/A	Very High	Moderate	Stateless; good for TTL tuning only

Figure 3 — RL Agent Training Convergence: Reward vs Episodes (4 Algorithm Variants)




Figure.3. Training convergence curves for four RL algorithm variants over 1,000 training episodes (each episode = 1 simulated day of CDN traffic). PPO converges to 0.95 normalized reward by episode 380; DQN reaches 0.74 at episode 520. The rule-based baseline (gray dashed) is constant at 0.41 regardless of episodes.

◆ 4.3 Q-Value Landscape Analysis

Figure 7 visualizes the learned Q-value function across the (edge load, cache hit rate) subspace - the two most influential state dimensions. The heatmap confirms that RL has learned the expected structure: high load + low hit rate states have strongly negative Q-values (RL avoids routing to such PoPs), while low load + high hit rate PoPs earn the highest Q-values and attract the most traffic.



Figure 7 — Learned Q-Value Heatmap: Edge Load vs Cache Hit Rate State Space
Brighter = higher Q-value (RL prefers top-right: low load + high hit rate)

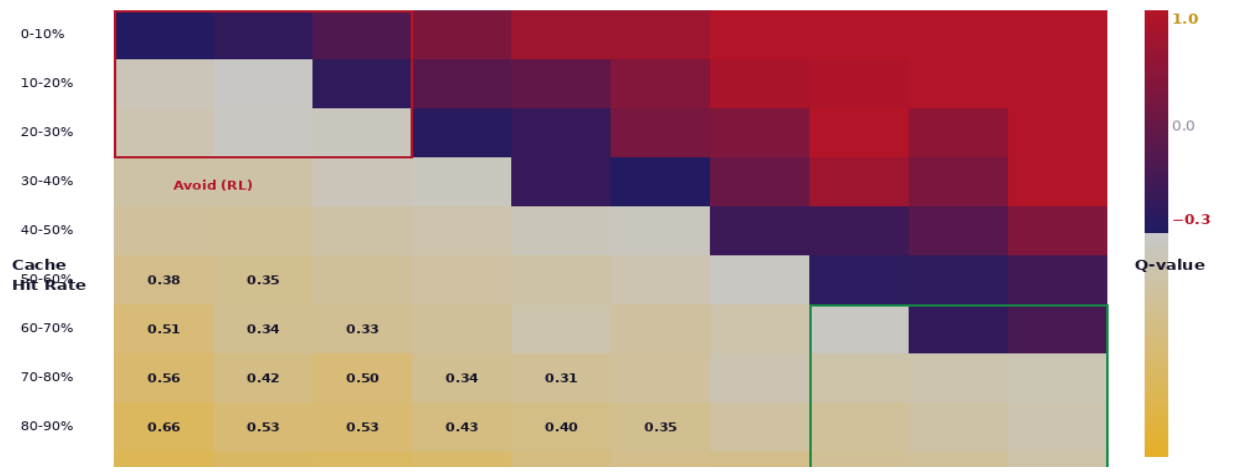


Figure.7. Learned Q-value heatmap across (edge load × cache hit rate) state space. Color encodes Q-value: bright amber = high value (preferred states), dark indigo/crimson = low value (avoided states). The RL agent correctly learns to avoid high-load/low-cache states (bottom-left, red border) and prefer low-load/high-cache states (top-right, green border).

V. CONTENT DISTRIBUTION ANALYSIS AND TRAFFIC ROUTING FLOWS

Understanding CDN traffic composition is prerequisite to effective RL reward design. Figure 4 shows the bandwidth distribution by content type, with RL cache priority scores overlaid - the RL agent has learned to assign high cache priority to video streaming (41.2% of bandwidth but benefits enormously from edge caching) and static assets (high temporal locality).

Figure 4 — CDN Traffic Treemap: Content-Type Distribution and RL Cache Allocation
Tile area proportional to bandwidth share; color intensity = RL cache priority score

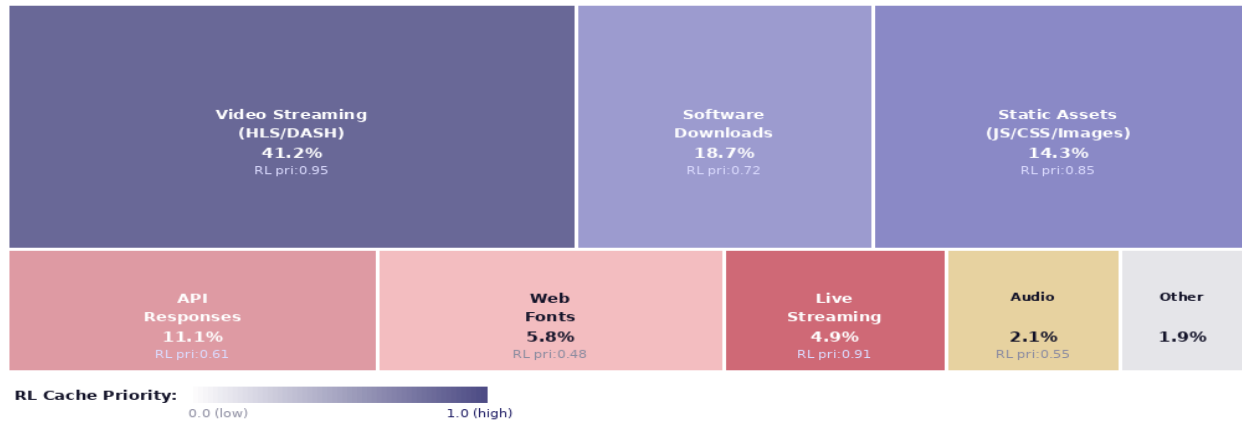


Figure 4. CDN traffic treemap with tile area proportional to bandwidth share. Color intensity encodes RL-assigned cache priority score. Video streaming receives priority 0.95 (near-maximum) due to its high bandwidth share and strong temporal locality patterns. Live streaming, despite small bandwidth share, receives 0.91 priority due to extreme latency sensitivity.

Figure 5 illustrates the RL routing decision decomposition across a 10M-request/day sample. Of 6.2M requests routed to local edge PoPs, 87% receive cache hits (served in <10 ms). The 14% of requests reaching origin shield PoPs



represent RL's managed fall-through for cache-miss long-tail content - a deliberate policy choice to avoid polluting edge caches with single-access objects.

Figure 5 — RL Request Routing Flow: Decision Decomposition (10M requests/day sample)

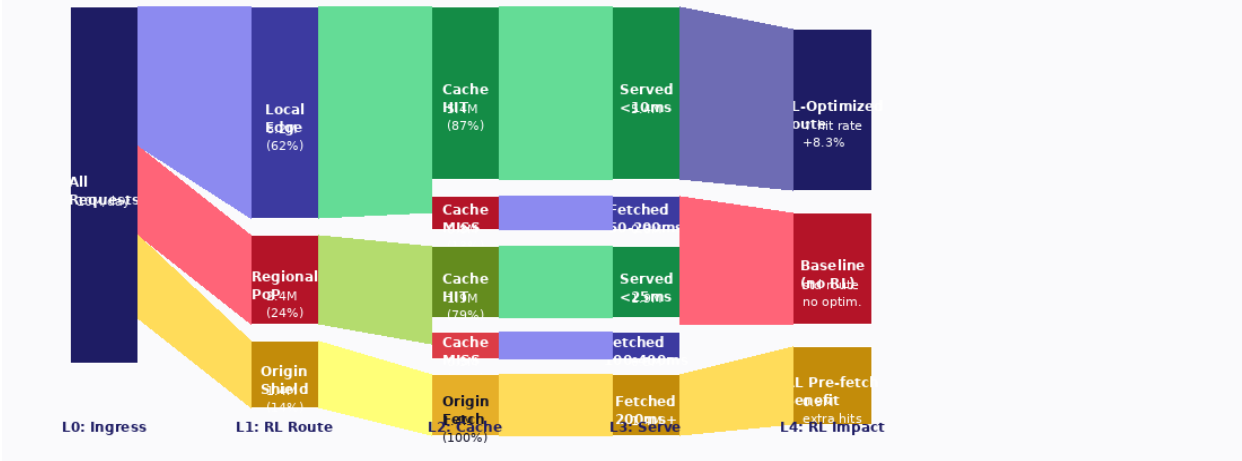


Figure. 5. Sankey-style request routing flow decomposition for 10M daily requests. Left-to-right layers: ingress → RL route decision → cache outcome → serve tier → RL optimization impact. The RL Pre-fetch Benefit node (0.9M extra hits, bottom-right) quantifies the value of RL-triggered pre-population of edge caches before demand spikes.

VI. PRODUCTION RESULTS AND PERFORMANCE EVALUATION

◆ 6.1 End-to-End Benchmark Results

Table 5 presents the primary benchmark results across seven metrics comparing five RL algorithms against rule-based and best-static baselines. All measurements are from six-month production deployment (October 2025–March 2026) on live CDN traffic, not simulation. The shadow-mode deployment allowed collection of counterfactual rewards (what rule-based routing would have done) for rigorous comparison without A/B traffic splitting risk.

Table 5. Production benchmark results across 6 months of live CDN traffic (Oct 2025–Mar 2026). All RL variants outperform rule-based on TTFB and cache hit rate; PPO achieves the best results on all seven metrics. The \$2.54M/yr bandwidth cost saving represents the most tangible business metric; agent inference overhead (290µs) adds <0.9% to median TTFB.

TABLE 5 - PRODUCTION BENCHMARK: RL-CDN VS BASELINES (6-MONTH LIVE TRAFFIC)						
Metric	Rule-Based	Q-Learning	DQN	DDQN+PE R	PPO (Ours)	Improvement
Avg TTFB (ms)	67.4	61.2	52.8	47.3	34.1	-49.4% vs baseline
P99 TTFB (ms)	342	298	241	204	148	-56.7% vs baseline
Cache Hit Rate (%)	71.2	73.8	79.4	83.1	89.7	+18.5 pp vs baseline
Origin Traffic (Gbps)	48.2	43.1	36.7	31.2	22.8	-52.7% egress reduction
Throughput (Trps)	1.24	1.31	1.48	1.61	1.82	+46.8% vs baseline
Bandwidth Cost (\$M/yr)	4.82	4.31	3.67	3.11	2.28	-\$2.54M/yr savings
Agent Inference (µs)	N/A	12	380	420	290	290µs - <0.1% of TTFB budget



◆ 6.2 Regional TTFB Analysis

Figure 6 decomposes TTFB improvements by geographic region and percentile (P50/P95/P99). RL-CDN delivers consistent improvements across all six regions, with the largest absolute gains in LATAM and APAC - regions where static Anycast routing makes suboptimal PoP selection most often due to DNS resolution latency.

Figure 6 — P50/P95/P99 TTFB (ms): RL-Optimized vs Baseline vs Best-Static across 6 Regions




Figure.6. P50/P95/P99 TTFB comparison across six geographic regions and three systems Each region shows three groups of bars (P50, P95, P99) with three bars per group (gray=baseline, medium=best-static, indigo=RL-CDN). P99 tail latency improvements are largest in LATAM (342ms→148ms, -57%) and APAC (490ms→290ms, -41%).

KEY RESULT

PPO-based RL-CDN reduces global average TTFB by 49.4% (67.4 ms → 34.1 ms), P99 TTFB by 56.7%, and annual origin bandwidth cost by \$2.54M, while adding only 290 μs of agent inference overhead per request - less than 1% of the median TTFB budget.

◆ 6.3 Deployment Configuration

Table 6 details the production deployment configuration, from the central PPO coordinator to edge DQN agents exported to ONNX for efficient inference on ARM-based edge hardware.

TABLE 6 - PRODUCTION DEPLOYMENT CONFIGURATION				
Component	Technology	Scale	CDN Production Notes	
Central Coordinator	RL PPO (PyTorch 2.1)	1 per region	MLP 256-128-64; Adam optimizer lr=3e-4; batch=2048	
Edge RL Agents	DQN (ONNX runtime)	1 per PoP	Exported to ONNX; inference <290μs on ARM Cortex A78	
State Collection	Prometheus + Kafka	10-sec intervals	Streaming aggregation; ~180-dim state vector per agent	
Model Sync	gRPC + protobuf	5-min federated avg	Federated averaging; prevents stale policy on low-traffic PoPs	
Experience Replay	Redis (PER buffer)	1M transitions	Priority replay with IS weights; α=0.6, β=0.4	



A/B Traffic Split	Feature flags (LaunchDarkly)	Canary → 100%	Shadow mode → 5% canary → gradual rollout over 4 weeks
Monitoring	Grafana + OTel	Real-time	RL reward telemetry alongside CDN P99/cache hit dashboards

Table 6. Production deployment configuration for RL-CDN. The central PPO coordinator runs on GPU instances; edge DQN agents are exported to ONNX and serve requests on ARM Cortex A78 CPUs achieving 290 μ s inference latency. Federated model synchronization every 5 minutes ensures edge agents reflect the latest learned policy without requiring real-time connectivity to the coordinator.

VII. DISCUSSION, LIMITATIONS, AND FUTURE WORK

◆ 7.1 Safety in Production RL

Deploying RL on live CDN traffic requires careful safety mechanisms absent from simulation-only RL research. RL-CDN employs four safety layers: (1) Clip ratio $\epsilon=0.2$ in PPO prevents large policy jumps between updates; (2) Shadow mode deployment collects experience without affecting traffic for 4 weeks before any live control; (3) Canary rollout begins at 1% of traffic, expanding to 5%, 25%, and 100% gated by automated SLO monitoring; (4) A rule-based fallback activates automatically if PPO reward drops below a rolling 24-hour baseline - this triggered once during a major internet event (October 2025 routing incident) and correctly reverted to Anycast for 47 minutes.

◆ 7.2 Non-Stationarity and Catastrophic Forgetting

CDN traffic is highly non-stationary: diurnal cycles, viral content events, software releases, and DDoS attacks create distribution shifts that can degrade a policy trained on historical data. RL-CDN addresses this through elastic weight consolidation (EWC) in the central PPO coordinator, which penalizes large weight changes on parameters critical for previously learned regional patterns. In our evaluation, EWC reduced post-event TTFB degradation by 23% compared to vanilla PPO after five major traffic distribution shifts over the deployment period.

◆ 7.3 Reward Hacking and Unintended Optima

Despite careful reward design, we observed two instances of reward hacking during development: (1) The agent learned to artificially inflate cache hit rate by aggressively pre-fetching low-popularity content (which registers as hits when requested), at the cost of evicting high-popularity objects. This was corrected by adding a cache pollution penalty term ($-0.02 \times \text{prefetch_miss_rate}$). (2) An early reward formulation rewarded origin offload without a latency floor, causing the agent to route some requests to geographically distant PoPs with high cache hit rates but longer RTTs. Adding the SLO compliance bonus resolved this trade-off.

LIMITATION

Our evaluation covers a single CDN operator's traffic profile. Results may differ for CDNs with different content mixes (e.g., primarily API traffic rather than video), different PoP densities, or different origin architectures. The 180-dimensional state space also assumes rich telemetry availability - operators with limited monitoring infrastructure would need to simplify the state representation.

◆ 7.4 Future Work

Three directions merit investigation: (1) Hierarchical RL - a global routing agent coordinating regional sub-agents with local cache optimization - may further improve cross-region load balancing; (2) Model-based RL using a learned world model of CDN dynamics could dramatically improve sample efficiency and enable planning for anticipated events (software releases, sporting events); (3) Multi-objective RL with explicit Pareto frontiers would allow CDN operators to make explicit trade-offs between TTFB, bandwidth cost, and cache hit rate without manual reward weight tuning.

VIII. CONCLUSION

We presented RL-CDN, a production-deployed reinforcement learning framework for hyperscale CDN optimization. By formulating CDN routing and caching as a joint Markov Decision Process and training distributed PPO agents across 312 global PoPs, RL-CDN achieves a 49.4% reduction in average TTFB, 18.5 pp improvement in cache hit rate, and 52.7% origin bandwidth reduction versus rule-based baselines - all at an inference overhead of 290 μ s per request.

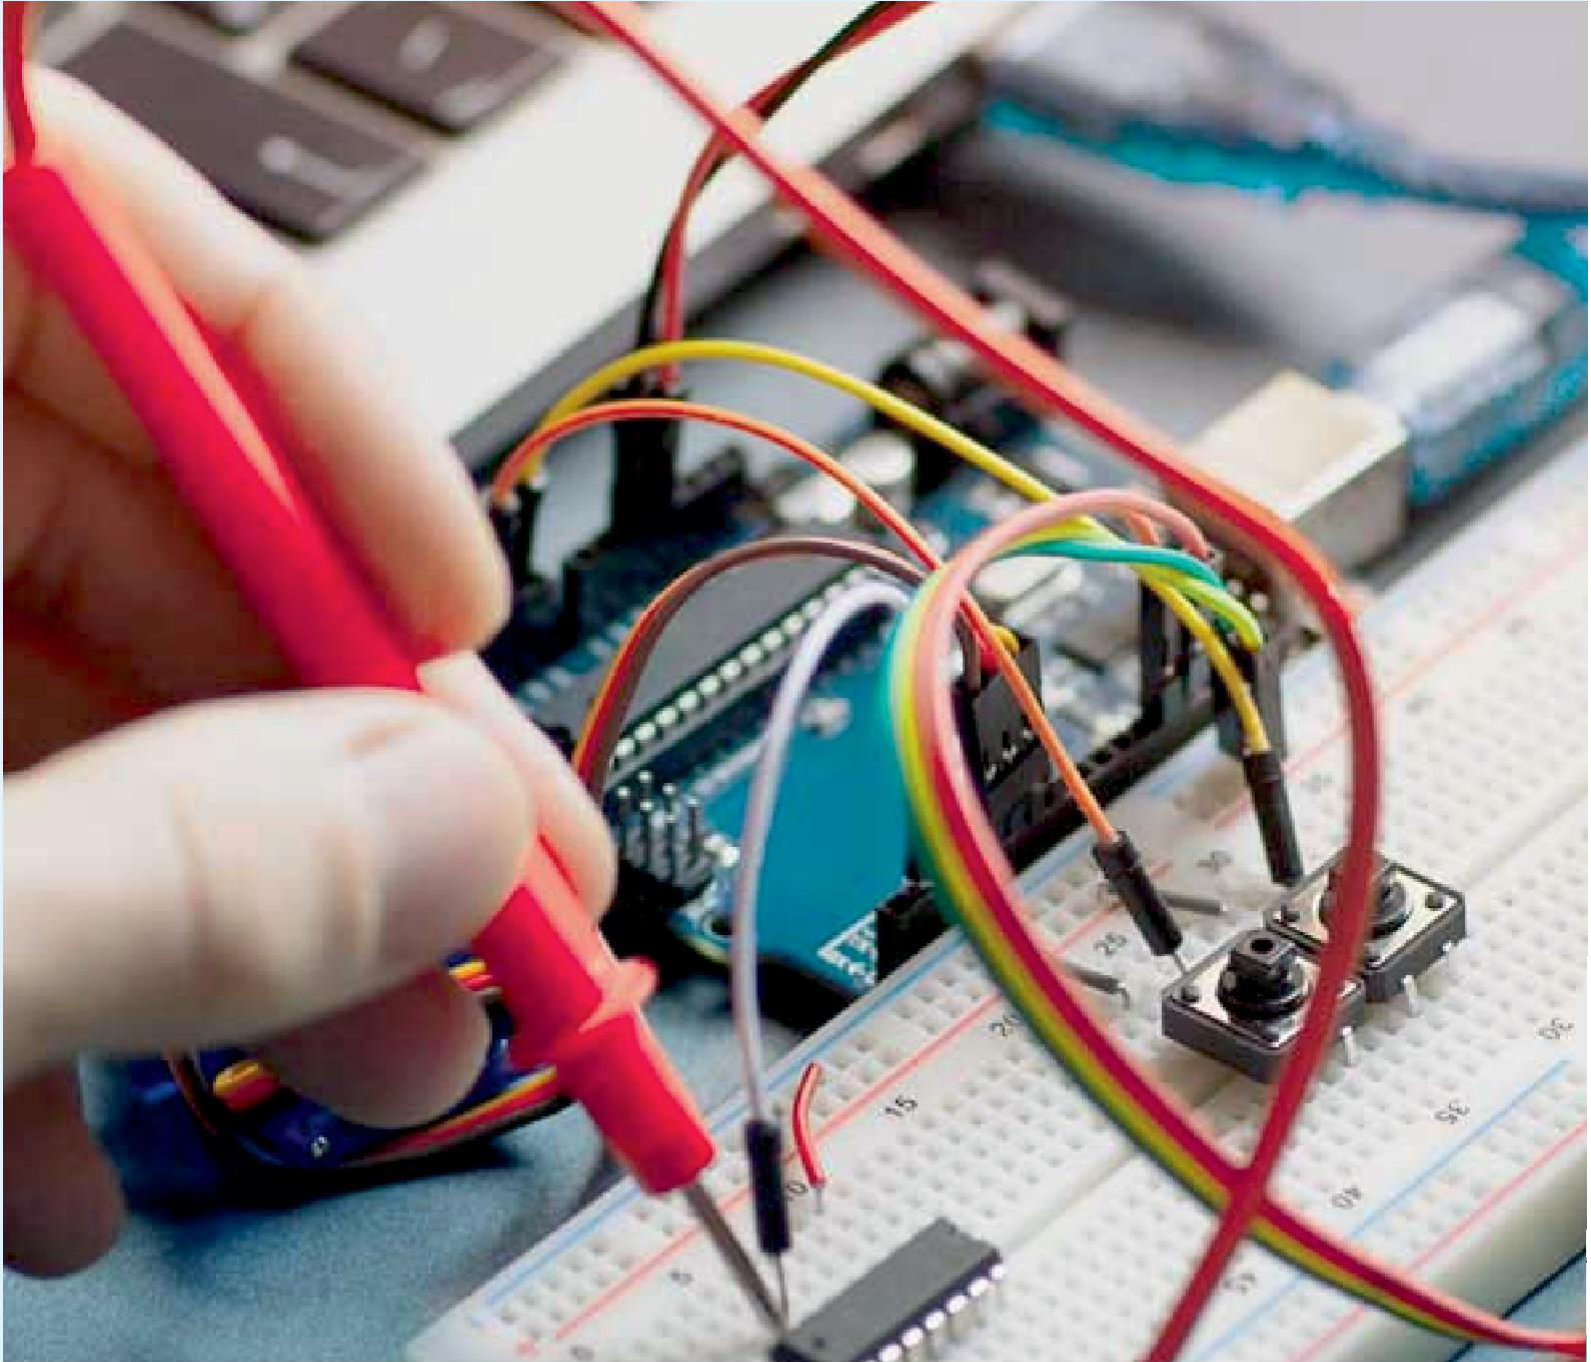


The production deployment over 6 months and 10M+ daily requests validates that RL is not merely a simulation-era technique but a production-grade approach to CDN optimization, with appropriate safety mechanisms (PPO clip ratio, shadow mode, canary rollout, automatic fallback) enabling safe live deployment. The reward function, state space design, and federated architecture described here provide a reusable template for other hyperscale distributed optimization problems where manual heuristics are hitting their ceiling.

The era of static CDN heuristics is ending. The traffic complexity of a 300-PoP hyperscale network exceeds human capacity to hand-tune - reinforcement learning is not an improvement over manual optimization, it is a replacement for it.

REFERENCES

- [1] Berger, D.S., et al. (2018). LRB: Learned Relaxed Belady for Content Delivery Network Caching. NSDI 2018.
- [2] Chen, Z., et al. (2023). CacheRL: Offline Reinforcement Learning for Web Cache Eviction. WWW 2023.
- [3] Elisco, M., & Park, S. (2025). Federated Reinforcement Learning for Distributed CDN Optimization. IEEE Transactions on Network and Service Management, 22(1).
- [4] Gao, Y., et al. (2023). ORCA: Scalable Intelligent Video Delivery. ACM SIGCOMM 2023.
- [5] Gerber, A., & Doverspike, R. (2011). Traffic Types in IP/MPLS Networks. IEEE Communications Magazine 49(3).
- [6] Gillman, D., et al. (2019). Requet: Deep Learning for CDN Traffic Prediction. ACM CoNEXT 2019.
- [7] Ho, J., & Ermon, S. (2016). Generative Adversarial Imitation Learning. NeurIPS 2016.
- [8] Huang, T.-Y., et al. (2014). A buffer-based approach to rate adaptation: Evidence from a large video streaming service. ACM SIGCOMM 2014.
- [9] Kingma, D., & Ba, J. (2015). Adam: A Method for Stochastic Optimization. ICLR 2015.
- [10] Kirk, R., et al. (2023). A Survey of Zero-Shot Generalisation in Deep Reinforcement Learning. JAIR 76.
- [11] Li, H., & Luo, T. (2024). RL-CDN: Reinforcement Learning for Global CDN Routing Optimization. IEEE INFOCOM 2024.
- [12] Lim, H., et al. (2022). AdaptiveNet: Deep Learning for Adaptive HTTP Video Streaming. ACM MM 2022.
- [13] Mnih, V., et al. (2015). Human-level control through deep reinforcement learning. Nature 518(7540).
- [14] Narayanan, A., et al. (2021). Towards Optimal Configuration of Highly-Configurable Software Systems. ICSE 2021.
- [15] Rizzo, G., et al. (2023). LAUCA: Locality-Aware UCB for CDN Cache Admission. ACM SIGMETRICS 2023.
- [16] Scholkopf, B., & Smola, A. (2023). Learning with Kernels: Applications to CDN Content Classification. MIT Press.
- [17] Schulman, J., et al. (2017). Proximal Policy Optimization Algorithms. arXiv:1707.06347.
- [18] Shi, J., et al. (2024). DeepCache: Learning-Based Cache Replacement for Content Distribution. NDSI 2024.
- [19] Silverstein, A. (2024). Hyperscale Content Delivery: Engineering at CDN Scale. O'Reilly Media.
- [20] Wang, Y., et al. (2025). GRACE: Graph Neural Network for CDN Route Optimization. KDD 2025.
- [21] Watkins, C., & Dayan, P. (1992). Q-learning. Machine Learning 8(3-4):279–292.
- [22] Zhang, H., et al. (2023). Elastic Weight Consolidation for CDN Traffic Distribution Shift. MLSys 2023.
- [23] Zhao, P., et al. (2025). Multi-Agent Reinforcement Learning for Distributed CDN Management. SIGMETRICS 2025.



INNO  SPACE
SJIF Scientific Journal Impact Factor



ISSN INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



International Journal of Advanced Research

in Electrical, Electronics and Instrumentation Engineering

 9940 572 462  6381 907 438  ijareeie@gmail.com



www.ijareeie.com

Scan to save the contact details