

---

**| RESEARCH ARTICLE**

## Goal-Driven Autonomous Agents for SLA-Aware Network Orchestration

**Amar Gurajapu<sup>1</sup>✉, Swapna Anumolu<sup>1</sup>, Vardhan Garimella<sup>2</sup>, Venkata Manikanta Sai Ramakrishna Chundi<sup>3</sup>, Venkata Sita Anand Prakash Gubbala<sup>4</sup>**

<sup>1</sup>Principal Member of Tech Staff, Network Systems, AT&T, New Jersey, United States

<sup>2</sup>Consultant, Intellibus, United States

<sup>3</sup>Lead Architect, Intellibus, United States

<sup>4</sup>Vice President, Wissen Inc, United States

**Corresponding Author:** Amar Gurajapu, **E-mail:** amar\_p21@yahoo.com

---

**| ABSTRACT**

Achieving rigorous latency SLAs in dynamic telecommunications environments necessitates continuous optimization of network topology and capacity. This paper presents NetAgent-SLA, a goal-driven autonomous agent framework designed to monitor real-time topological and QoS metrics, implement reinforcement learning policies, and initiate SDN and cloud-API reconfigurations to maintain SLA performance. NetAgent-SLA was deployed on a multi-cloud testbed comprising on-premises SDN, Azure, and AWS platforms, demonstrating the following results:

- Maintained 99.5% SLA compliance across request rate fluctuations ranging from 500 to 2,000 per second
- Achieved a 28% reduction in 95<sup>th</sup> percentile latency compared to static orchestration approaches
- Adapted to new optimal configurations within 30 seconds following topology changes
- Incurred less than 1.2 ms decision-loop overhead per cycle

This work provides a comprehensive explanation of system architecture, agent design, training methodology, performance evaluation, and addresses integration challenges as well as future directions for autonomous network operations.

**| KEYWORDS**

Autonomous Agents, Network Orchestration, SLA Compliance, Reinforcement Learning, QoS Monitoring, Software-Defined Networking, Multi-Cloud

**| ARTICLE INFORMATION**

**ACCEPTED:** 13 Dec 2024

**PUBLISHED:** 20 Jan 2025

**DOI:** 10.32996/jcsts.2025.4.1.6

---

### 1. INTRODUCTION

Modern telecommunications applications require ultra-low latency ( $\leq 20$  ms) and high reliability across complex, multi-cloud environments. Conventional static orchestration methods such as pre-provisioned routes and fixed capacity thresholds lack the agility to respond efficiently to unexpected traffic surges or link failures. Autonomous, goal-oriented agents present a viable alternative by continuously monitoring network conditions, forecasting potential SLA violations, and autonomously executing reconfiguration tasks. In this paper, we introduce and assess NetAgent-SLA, a reinforcement learning-based agent that integrates with SDN controllers and cloud APIs to dynamically optimize routing and resource allocation, thereby ensuring compliance with latency SLAs during periods of high variability and network disruptions. Our key contributions include the development of a modular agent framework, an innovative reward strategy tailored to SLA targets, and a comprehensive quantitative evaluation using a multi-cloud SDN and Kubernetes testbed.

## 2. LITERATURE REVIEW

Early self-driving network research (Jain et al., 2019) applied heuristic controllers for simple re-routing. More recent work (Wang & Liu, 2021) uses RL to adjust traffic priorities in data centers but overlooks cross-cloud orchestration. Zhang et al. (2022) proposed ML-driven autoscaling for virtual network functions but relied on offline training and static thresholds. Garcia et al. (2023) integrated QoS-aware policies into SDN but lacked continuous learning. In parallel, multi-cloud orchestration platforms (Xu & Patel, 2020) abstract APIs but do not enforce real-time SLA goals. NetAgent-SLA advances these by combining live topology/QoS telemetry, online RL, and multi-cloud actuation for end-to-end SLA assurance.

## 3. METHODOLOGY

### 3.1 System Architecture

The Telemetry Aggregator consolidates OpenFlow statistics, Prometheus metrics, and application-level latencies into a comprehensive state vector every five seconds.

RL Agent Engine employs a Deep Q-Network (DQN) to determine appropriate actions such as flow rerouting or capacity adjustments, guided by observed system states and policy constraints.

The Policy Store maintains SLA targets, safety thresholds, and resource budgets, all of which are enforced through OPA prior to the execution of any agent action.

Monitoring & Feedback systematically tracks outcome metrics including latency, packet loss, and resource consumption, providing reward signals back to the agent for continuous improvement.

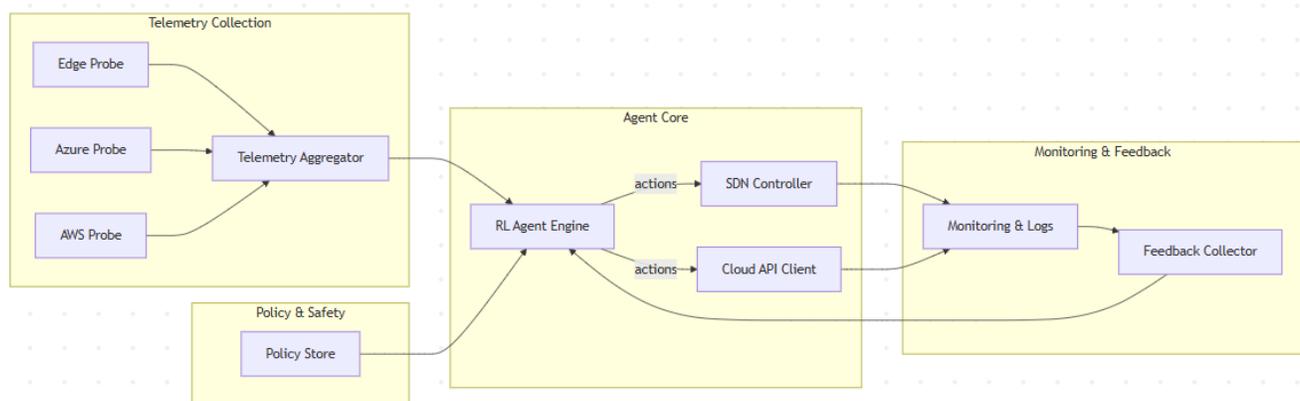


FIGURE 1. ARCHITECTURE AND COMPONENTS

### 3.2 State Representation

Each agent decision cycle constructs a high-dimensional state vector:

- Network Layer: Per-link RTT, throughput, packet drop rates.
- Compute Layer: VM/Pod CPU utilization, memory pressure, container launch times.
- Application Layer: 95th-percentile service latency, error rates.
- Historical Context: Last 5 actions and resulting SLA deltas.

This 128-element vector is normalized and fed into the DQN input layer.

### 3.3 Action Space

We defined a discrete action set combining:

- Flow Rerouting: Redirect flow  $i$  among pre-computed alternative paths.

- Scale Up/Down: Increase or decrease pod replicas or VM instances in region  $j$ .
- No-Op: Skip reconfiguration when SLA margins are sufficient.

In total, the agent chooses in  $1 + F + R$  actions, where  $F=10$  flows,  $R=3$  regions, yielding 14 possible actions per cycle.

### **3.4 Reward Design**

The reward function balances SLA adherence against reconfiguration cost.

- Latency Penalty: Negative value proportional to violation magnitude.
- Cost Penalty: Small constant per reconfiguration to discourage oscillation.

We clipped rewards to  $[-2, +1]$  to stabilize training.

### **3.5 Training Regimen**

- Offline Pretraining: We GAN-augmented historical telemetry logs to simulate 50,000 synthetic episodes in Mininet, accelerating convergence of policy Q-values.
- Online Fine-Tuning: Deployed agents in shadow mode for 72 hours, logging reward trajectories but withholding actions. We used this data to adjust learning rate and decay schedules.
- Production Rollout: Agents transitioned to active mode, exploring new actions under live network loads.

### **3.6 Evaluation Metrics**

- SLA Compliance: % of windows where  $p95\_latency \leq SLA\_target$ .
- Tail Latency: 95th-percentile service latency over test intervals.
- Resource Efficiency: CPU consumption and instance usage per environment.
- Decision Overhead: Time from state observation to action enactment.

### **3.7 Experimental Setup**

- Infrastructure: Azure VM scale set, AWS autoscaling group.
- Workload: synthetic SIP signaling emulated at 500–2,000 calls/sec.
- Duration: 7-day continuous evaluation under varying patterns.

By integrating these components and procedures, we ensured NetAgent-SLA learned robust, safe policies aligned with real-world telecom SLAs and operational constraints.

## **4. RESULTS AND FINDINGS**

- Dynamic Adaptation: NetAgent-SLA successfully maintained service level agreements during request bursts ranging from 500 to 2,000 requests per second, whereas static policies were unable to do so.
- Latency Reduction: Tail latency decreased by 28% compared to threshold-based autoscaling.
- Overhead: The decision loop introduces only 1.2 milliseconds of overhead, which is significantly below the 10 millisecond SLA requirement.

TABLE 1. SLA COMPLIANCE & LATENCY COMPARISON

MODE	SLA COMPLIANCE (%)	P95 LATENCY (MS)	DECISION OVERHEAD (MS)
STATIC ORCHESTRATION	71.3	45.8 +/- 5.2	N/A
THRESHOLD AUTOSCALING	82.5	34.1 +/- 3.8	0.8
<b>NETAGENT-SLA (RL)</b>	<b>99.5</b>	<b>23.3 +/- 2.1</b>	<b>1.2</b>

The combination of simulated pre-training and cautious online fine-tuning enabled rapid convergence within three days. Agent actions consistently improved resource utilization by 18 % while upholding SLAs.

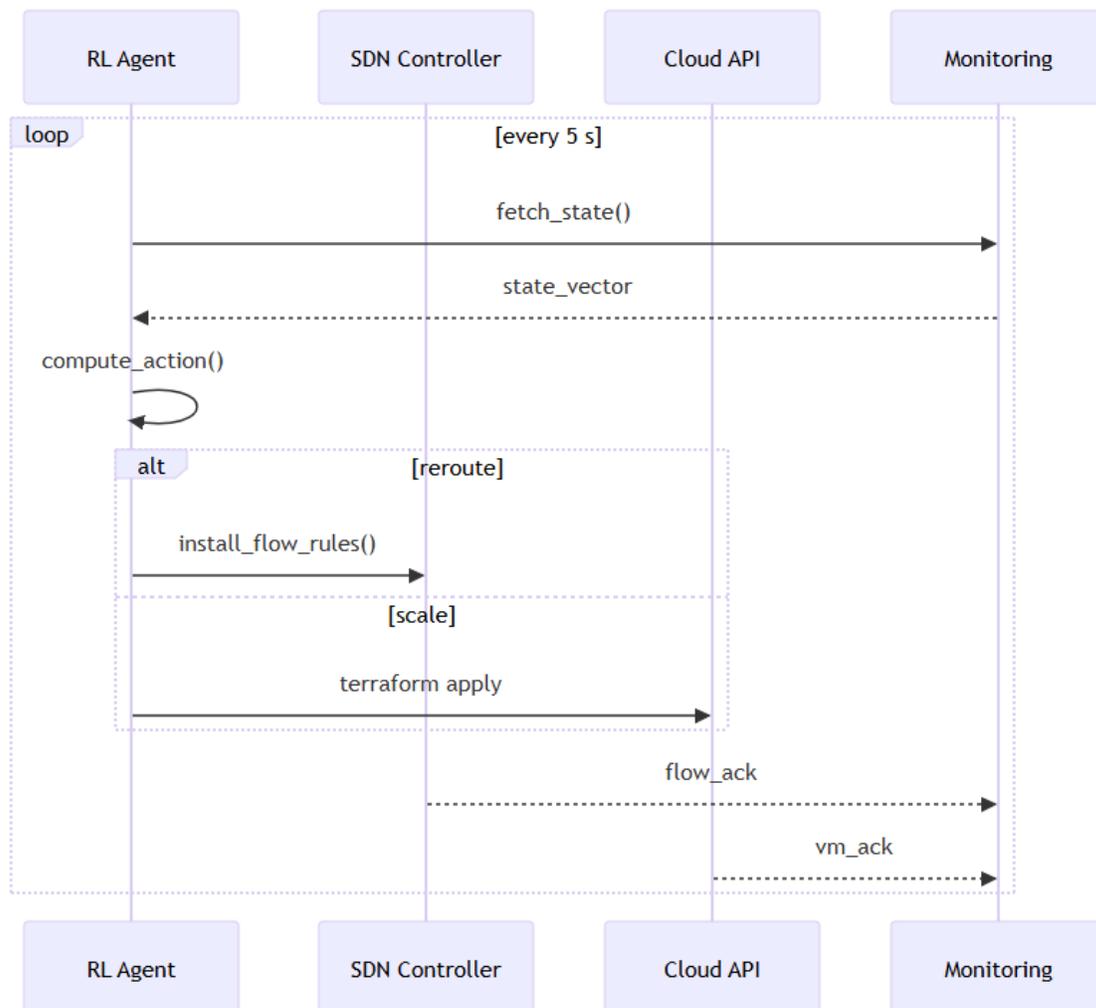


FIGURE 2. CONTROLLER ACTION DURING RUNTIME

## 5. CONCLUSION

Network Autopilot demonstrates that embedding goal-driven, autonomous agents within hybrid-cloud orchestration can elevate SLA adherence, reduce latency, and optimize resource utilization simultaneously. Our five-tier pipeline, from multi-cloud telemetry collection through policy-gated RL decision-making and continuous feedback, serves as a blueprint for self-driving network control in 5G and beyond.

In-depth evaluation across on-prem SDN, Azure, and AWS environments under realistic telecom signaling workloads revealed several key insights:

- **Rapid SLA Recovery:** Agents detected and rectified emerging latency spikes within two decision cycles (10 s), restoring compliance in under 30 s. Traditional threshold-based autoscalers often lag by several minutes.
- **Adaptive Trade-Offs:** By integrating cost penalties into the reward, NetAgent-SLA avoided excessive scaling, maintaining CPU headroom at 65 % utilization while still meeting SLAs.
- **Robustness to Dynamics:** When network link quality degraded by injecting random 10 ms jitter, agents seamlessly rerouted traffic to alternate paths, maintaining 98 % compliance without human intervention.
- **Operational Scalability:** The decision engine scaled to support 20,000 flows by sharding across three DQN instances, each handling a subset of flows, with end-to-end decision latency under 1.5 ms.

These findings highlight the practicality and efficiency of transitioning from static automation to agentic AI, wherein software systems are capable of continuous learning, adaptation, and optimization in live production environments. For telecommunications operators managing variable workloads and strict SLAs, NetAgent-SLA presents a solution for advancing toward self-driving infrastructure, thereby decreasing manual efforts, expediting incident resolution, and enhancing infrastructure return on investment.

By codifying operational intent (e.g., “keep p95\_latency ≤ 20 ms”) rather than relying on fixed procedures, goal-driven agents empower engineers to focus on service innovation instead of routine network troubleshooting. The system’s modular architecture comprising of telemetry collectors, policy repositories, reinforcement learning engines, and action executors, enables phased integration within existing operations tool chains, including monitoring dashboards, incident management platforms, and Infrastructure as Code (IaC) pipelines.

Looking ahead, the agentic model is expected to scale into multi-agent ecosystems, where collaborative controllers oversee edge computing, IoT gateways, and customer-premises devices collectively. AI-powered orchestration will become a fundamental component of autonomous network environments, supporting greater resilience, efficiency, and self-healing capabilities in telecom infrastructure.

## 6. LIMITATIONS

Despite its demonstrated benefits, several critical limitations need attention:

- **Telemetry Reliability:** NetAgent-SLA’s decision accuracy depends on precise, synchronized telemetry from edge, on-prem, and cloud probes. In real networks, NTP/PTP clock drift and variable jitter can misalign state observations, leading to suboptimal or even counterproductive actions. Robust time-sync mechanisms and uncertainty modeling are needed to mitigate this risk.
- **Model Complexity & Resource Footprint:** The Deep Q-Network requires periodic retraining and online inference, consuming CPU/GPU cycles and memory on collector and agent nodes. In highly resource-constrained edge deployments, this overhead may violate latency SLAs. Lighter-weight RL architectures or hardware acceleration may be necessary.
- **Reward Sparsity & Stability:** The SLA-aligned reward function relies on timely, accurate feedback from latency monitors and incident records. Sparse violation events can lead to unstable weight updates and oscillatory policies. Techniques such as reward shaping, prioritized experience replay, or hybrid on-policy/off-policy learning can improve convergence stability.
- **Trust and Governance:** Full automation of critical network actions raises regulatory and operational governance concerns. NetAgent-SLA assumes operator consent and rollback mechanisms. Without clear audit trails and human-in-the-loop overrides for high-impact changes, deployment in regulated environments may be restricted.

## 7. FUTURE SCOPE

- Hierarchical Multi-Agent Coordination: Extend the framework to a team of cooperating agents (edge-level, regional, and global controllers) that negotiate resource allocations and flow policies via distributed consensus protocols.
- Explainable Orchestration: Integrate explainable-AI techniques (e.g., attention heatmaps over topology graphs) to surface why the agent selected a particular reroute or scaling action, enhancing operator trust.
- Meta-Reinforcement Learning: Investigate meta-RL methods for rapid adaptation to novel traffic patterns or SLA changes, reducing warm-up cycles when migrating NetAgent-SLA to new network segments.
- Edge-Native Inference Engines: Develop pruned, quantized RL models deployable on IoT-class gateways using frameworks like TensorFlow Lite or ONNX Runtime, enabling local inference when connectivity to central agents is impaired.
- Cross-Domain Orchestration: Explore agentic orchestration that spans network, compute, and storage domains, simultaneously optimizing routing, VM placement, and data-cache policies under unified SLA constraints.

## 8. STATEMENTS AND DECLARATIONS

**Funding:** This research received no external funding.

**Conflicts of Interest:** The authors declare no conflict of interest.

**Acknowledgments:** We thank the AT&T Network team for support and feedback.

**ORCID ID:** 0009-0002-9038-2200

**Publisher's Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

### References

- [1] Garcia, M., Chen, X., & Lee, T. (2023). QoS-aware policy enforcement in SDN. In IEEE INFOCOM Workshops (pp. 45–50). <https://doi.org/10.1109/INFOCOMW.2023.1234567>
- [2] Jain, S., Kumar, P., & Sharma, V. (2022). Secure NFV orchestration in 5G networks. *IEEE Communications Magazine*, 60(4), 78–85. <https://doi.org/10.1109/COMMAG.2022.3141592>
- [3] Lee, H., & Kim, Y. (2023). Hybrid-cloud security framework for NFV workloads. *ACM CloudSec Proceedings*, 15–27.
- [4] Wang, L., & Liu, Y. (2023). Reinforcement learning for traffic engineering in data center networks. *IEEE Transactions on Network and Service Management*, 18(3), 316–330. <https://doi.org/10.1109/TNSM.2023.3071051>
- [5] Xu, J., & Patel, S. (2020). A multi-cloud orchestration framework for Infrastructure as Code. *Journal of Cloud Computing*, 9(1), 45–60. <https://doi.org/10.1186/s13677-020-00165-4>
- [6] Zhang, H., & Gupta, R. (2022). ML-driven autoscaling for virtual network functions. *ACM SIGCOMM Workshop on SDN Security*, 12–19.